

Third Edition

Robert  
Audi

# EPISTEMOLOGY

A Contemporary  
Introduction to the  
Theory of Knowledge

# Epistemology

This comprehensive introduction explains the concepts and theories central to understanding knowledge. The third edition features new sections on such topics as the nature of intuition, the skeptical challenge of rational disagreement, and “the value problem”—the question why knowledge is preferable to mere true belief. Special features of the third edition of *Epistemology* include:

- enhanced treatment of key topics such as perception, scientific hypotheses, self-evidence and the a priori, testimony, contextualism, understanding, and virtue epistemology
- expanded discussion of the relation between epistemology and related fields, especially philosophy of mind, philosophy of science, and ethics
- greater clarity for undergraduate readers.

## **Praise for the third edition:**

“ . . . Robert Audi’s *Epistemology*, Third Edition, is the most authoritative, comprehensive, and state-of-the-art textbook in the field. In clear, masterful prose, Audi covers all the main topics in epistemology . . . Every student of epistemology—new and old—should read this book.”

—Peter Graham, University of California, Riverside

“ . . . unusually comprehensive, elegantly structured, and accessible . . . a cutting-edge treatment of the latest debates about the nature of intuitions, the significance of rational disagreement, and the value of knowledge and justified true belief.”

—Ralph Kennedy, Wake Forest University

“ . . . a well-motivated, comprehensive, accessible introduction for students as well as an original, exciting, cutting-edge work of epistemology in its own right. Novices and experts alike will continually profit—and tremendously so—from studying it . . . an ideal text for undergraduate courses in epistemology, and even graduate-level surveys.”

—E.J. Coffman, University of Tennessee

### **Praise for the second edition:**

“... philosophically insightful and masterfully written—even more so in its new edition. Guaranteed to fascinate the beginner while retaining its exalted status with the experts.”

—Claudio de Almeida, PUCRS, Brazil

“My students like this book and have learned much from it, as I have . . . *Epistemology* . . . is simply the best textbook in epistemology that I know of.”

—Thomas Vinci, Dalhousie University

### **Praise for the first edition:**

“No less than one would expect from a first-rate epistemologist who is also a master expositor . . . A superb introduction.”

—Ernest Sosa, Rutgers University

“This is a massively impressive book, introducing . . . virtually all the main areas of epistemology. . . . lucid and highly readable, while not shirking the considerable complexities of his subject matter.”

—Elizabeth .M. Fricker, University of Oxford

“A state-of-the-art introduction to epistemology by one of the leading figures in the field.”

—William P. Alston, Syracuse University

**Robert Audi** is John A. O'Brien Professor of Philosophy at the University of Notre Dame and author of many papers and books on knowledge and belief, justification, and rationality.

# **Epistemology**

A Contemporary Introduction to the  
Theory of Knowledge

Third Edition

Robert Audi



**Routledge**

Taylor & Francis Group

NEW YORK AND LONDON



**Routledge Contemporary Introductions to Philosophy**  
Series editor: *Paul K Moser, Loyola University of Chicago*

This innovative, well-structured series is for students who have already done an introductory course in philosophy. Each book introduces a core general subject in contemporary philosophy and offers students an accessible but substantial transition from introductory to higher-level college work in that subject. The series is accessible to non-specialists and each book clearly motivates and expounds the problems and positions introduced. An orientating chapter briefly introduces its topic and reminds readers of any crucial material they need to have retained from a typical introductory course. Considerable attention is given to explaining the central philosophical problems of a subject and the main competing solutions and arguments for those solutions. The primary aim is to educate students in the main problems, positions, and arguments of contemporary philosophy rather than to convince students of a single position.

**Epistemology**

Third edition

*Robert Audi*

**Philosophy of Mind**

Third edition

*John Heil*

**Philosophy of Perception**

*William Fish*

**Metaphysics**

Third edition

*Michael J. Loux*

**Philosophy of Mathematics**

Second edition

*James Robert Brown*

**Classical Philosophy**

*Christopher Shields*

**Classical Modern**

**Philosophy**

*Jeffrey Tlumak*

**Ethics**

*Harry Gensler*

**Philosophy of Science**

Second edition

*Alex Rosenberg*

**Philosophy of Biology**

*Alex Rosenberg and Daniel W. McShea*

**Philosophy of Language**

Second edition

*William G. Lycan*

**Philosophy of Religion**

*Keith E. Yandell*

**Philosophy of Art**

*Noël Carroll*

**Philosophy of Psychology**

*José Bermudez*

**Social and Political**

**Philosophy**

*John Christman*

**Continental Philosophy**

*Andrew Cutrofello*

**Free Will (forthcoming)**

*Michael McKenna*

**Ethics (forthcoming)**

Second edition

*Harry Gensler*

**Feminist Philosophy  
(forthcoming)**

*Lorraine Code*

**Philosophy of Literature  
(forthcoming)**

*John Gibson*

First edition published 1998, second edition published 2003

by Routledge

This edition first published 2011

by Routledge

270 Madison Ave, New York, NY 10016

Simultaneously published in the UK

by Routledge

2 Park Square, Milton Park, Abingdon, Oxon OX14 4RN

*Routledge is an imprint of the Taylor & Francis Group, an informa business*

This edition published in the Taylor & Francis e-Library, 2010.

To purchase your own copy of this or any of Taylor & Francis or Routledge's collection of thousands of eBooks please go to [www.eBookstore.tandf.co.uk](http://www.eBookstore.tandf.co.uk).

First edition © 1998 Taylor & Francis

Second edition © 2003 Taylor & Francis

Third edition © 2011 Taylor & Francis

All rights reserved. No part of this book may be reprinted or reproduced or utilized in any form or by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying and recording, or in any information storage or retrieval system, without permission in writing from the publishers.

**Trademark Notice:** Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe.

*Library of Congress Cataloging in Publication Data*

Audi, Robert

Epistemology : a contemporary introduction to the theory of knowledge / Robert Audi. — 3rd ed.

p. cm. — (Routledge contemporary introductions to philosophy)

Includes bibliographical references and index.

[etc.]

I. Knowledge, Theory of. I. Title.

BD161.A783 2010

121—dc22

2010005139

ISBN 13: 978-0-415-87922-4 (hbk)

ISBN 13: 978-0-415-87923-1 (pbk)

To Malou





# Contents

<i>Preface to the first edition</i>	xiii
<i>Acknowledgments to the first edition</i>	xvi
<i>Preface to the second edition</i>	xviii
<i>Acknowledgments to the second edition</i>	xx
<i>Preface to the third edition</i>	xxi
<i>Acknowledgments to the third edition</i>	xxii
<b>Introduction: a sketch of the sources and nature of belief, justification, and knowledge</b>	<b>1</b>
<i>Perception, belief, and justification</i>	1
<i>Justification as process, as status, and as property</i>	2
<i>Knowledge and justification</i>	4
<i>Memory, introspection, and self-consciousness</i>	5
<i>Reason and rational reflection</i>	5
<i>Testimony</i>	6
<i>Basic sources of belief, justification, and knowledge</i>	6
<i>Three kinds of grounds of belief</i>	7
<i>Fallibility and skepticism</i>	8
<i>Overview</i>	9
<b>Part One</b>	
<b>Sources of justification, knowledge, and truth</b>	<b>13</b>
<b>I Perception: sensing, believing, and knowing</b>	<b>16</b>
<i>The elements and basic kinds of perception</i>	17
<i>Seeing and believing</i>	22
<i>Perceptual justification and perceptual knowledge</i>	26
<i>Notes</i>	31

<b>2</b>	<b>Theories of perception: sense experience, appearances, and reality</b>	38
	<i>Some commonsense views of perception</i>	38
	<i>The theory of appearing</i>	40
	<i>Sense-datum theories of perception</i>	41
	<i>Adverbial theories of perception</i>	47
	<i>Adverbial and sense-datum theories of sensory experience</i>	49
	<i>Phenomenalism</i>	51
	<i>Perception and the senses</i>	55
	<i>Notes</i>	59
<b>3</b>	<b>Memory: the preservation and reconstruction of the past</b>	62
	<i>Memory and the past</i>	63
	<i>The causal basis of memory beliefs</i>	64
	<i>Theories of memory</i>	66
	<i>Remembering, recalling, and imaging</i>	72
	<i>Remembering, imaging, and recognition</i>	74
	<i>The epistemological centrality of memory</i>	75
	<i>Notes</i>	79
<b>4</b>	<b>Consciousness: the life of the mind</b>	84
	<i>Two basic kinds of mental properties</i>	84
	<i>Introspection and inward vision</i>	86
	<i>Some theories of introspective consciousness</i>	87
	<i>Consciousness and privileged access</i>	91
	<i>Introspective consciousness as a source of justification and knowledge</i>	96
	<i>Notes</i>	100
<b>5</b>	<b>Reason I: understanding, insight, and intellectual power</b>	104
	<i>Self-evident truths of reason</i>	104
	<i>The classical view of the truths of reason</i>	107
	<i>The empiricist view of the truths of reason</i>	116
	<i>Notes</i>	121
<b>6</b>	<b>Reason II: meaning, necessity, and provability</b>	130
	<i>The conventionalist view of the truths of reason</i>	130
	<i>Some difficulties and strengths of the classical view</i>	133
	<i>Reason, experience, and a priori justification</i>	138
	<i>Notes</i>	145
<b>7</b>	<b>Testimony: the social foundation of knowledge</b>	150
	<i>The nature of testimony: formal and informal</i>	150
	<i>The psychology of testimony</i>	151

<i>The epistemology of testimony</i>	155
<i>The indispensability of testimonial grounds</i>	161
<i>Notes</i>	167

## Part Two

### **The structure and growth of justification and knowledge 173**

<b>8 Inference and the extension of knowledge</b>	176
<i>The process, content, and structure of inference</i>	177
<i>Inference and the growth of knowledge</i>	182
<i>Source conditions and transmission conditions for inferential knowledge and justification</i>	185
<i>Memorial preservation of inferential justification and inferential knowledge</i>	197
<i>Notes</i>	199
<b>9 The architecture of knowledge</b>	206
<i>Inferential chains and the structure of belief</i>	206
<i>The epistemic regress problem</i>	210
<i>The epistemic regress argument</i>	215
<i>Foundationalism and coherentism</i>	216
<i>Holistic coherentism</i>	217
<i>The nature of coherence</i>	221
<i>Coherence and second-order justification</i>	229
<i>Moderate foundationalism</i>	232
<i>Notes</i>	236

## Part Three

### **The nature and scope of justification and knowledge 243**

<b>10 The analysis of knowledge: justification, certainty, and reliability</b>	246
<i>Knowledge and justified true belief</i>	246
<i>Knowledge conceived as the right kind of justified true belief</i>	248
<i>Naturalistic accounts of the concept of knowledge</i>	253
<i>Problems for reliability theories</i>	257
<i>Notes</i>	264
<b>11 Knowledge, justification, and truth: internalism, externalism, and intellectual virtue</b>	270
<i>Knowledge and justification</i>	270
<i>Internalism and externalism in epistemology</i>	272
<i>Internalist and externalist versions of virtue epistemology</i>	277

	<i>Justification, knowledge, and truth</i>	281
	<i>The value problem</i>	282
	<i>Theories of truth</i>	286
	<i>Concluding proposals</i>	290
	<i>Notes</i>	292
<b>12</b>	<b>Scientific, moral, and religious knowledge</b>	298
	<i>Scientific knowledge</i>	298
	<i>Moral knowledge</i>	308
	<i>Religious knowledge</i>	319
	<i>Notes</i>	328
<b>13</b>	<b>Skepticism I: the quest for certainty</b>	334
	<i>The possibility of pervasive error</i>	334
	<i>Skepticism generalized</i>	337
	<i>The egocentric predicament</i>	343
	<i>Fallibility</i>	344
	<i>Uncertainty</i>	347
	<i>Notes</i>	353
<b>14</b>	<b>Skepticism II: the defense of common sense in the face of fallibility</b>	358
	<i>Negative versus positive defenses of common sense</i>	358
	<i>Deducibility, evidential transmission, and induction</i>	359
	<i>The authority of knowledge and the cogency of its grounds</i>	362
	<i>Refutation and rebuttal</i>	365
	<i>Prospects for a positive defense of common sense</i>	366
	<i>The challenge of rational disagreement</i>	371
	<i>Skepticism and common sense</i>	374
	<i>Notes</i>	375
<b>15</b>	<b>Conclusion</b>	379
	<i>Short annotated bibliography of books in epistemology</i>	389
	<i>Index</i>	399

# **Preface to the first edition**

This book is a wide-ranging introduction to epistemology, conceived as the theory of knowledge and justification. It presupposes no special background in philosophy and is meant to be fully understandable to any generally educated, careful reader, but for students it is most appropriately studied after completing at least one more general course in philosophy.

The main focus is the body of concepts, theories, and problems central in understanding knowledge and justification. Historically, justification—sometimes under such names as ‘reason to believe’, ‘evidence’, and ‘warrant’—has been as important in epistemology as knowledge itself. This is surely so at present. In many parts of the book, justification and knowledge are discussed separately; but they are also interconnected at many points. The book is not historically organized, but it does discuss selected major positions in the history of philosophy, particularly some of those that have greatly influenced human thought. Moreover, even where major philosophers are not mentioned, I try to take their views into account. One of my primary aims is to facilitate the reading of those philosophers, especially their epistemological writings. It would take a very long book to discuss representative contemporary epistemologists or, in any detail, even a few historically important epistemologies, but a shorter one can provide many of the tools needed to understand them. Providing such tools is one of my main purposes.

The use of this book in the study of philosophy is not limited to courses or investigations in epistemology. Epistemological problems and theories are often interconnected with problems and theories in the philosophy of mind; nor are these two fields of philosophy easily separated (a point that may hold, if to a lesser extent, for any two central philosophical fields). There is, then, much discussion of the topics in the philosophy of mind that are crucial for epistemology, for instance the phenomenology of perception, the nature of belief, the role of imagery in memory and introspection, the variety of mental properties figuring in self-knowledge, the nature of inference, and the structure of a person’s system of beliefs.

Parts of the book might serve as collateral reading not only in pursuing the philosophy of mind but also in the study of a number of philosophers often discussed in philosophy courses, especially Plato, Aristotle, Aquinas,

Descartes, Leibniz, Locke, Berkeley, Hume, Kant, and Mill. The book might facilitate the study of moral philosophy, such as Kantian and utilitarian ethics, both discussed in some detail in Chapter 9; and it bears directly on topics in the epistemology of religion, some of which are also discussed in Chapter 9.

The writing is intended to be as simple and concrete as possible for a philosophically serious introduction that does not seek simplicity at the cost of falsehood. The territory surveyed, however, is extensive and rich. This means that the book cannot be traversed quickly without missing landmarks or failing to get a view of the larger segments and their place in the whole. Any one chapter can perhaps be read at a sitting, but experience has shown that even the shortest chapter covers too many concepts and positions for most readers to assimilate in a single reading and far more than most instructors can cover in any detail in a single session.

To aid concentration on the main points, and to keep the book from becoming more complicated, notes are limited, though parenthetical references are given in some places and there is also a short selected bibliography with thumbnail annotations. By and large, the notes are not needed for full comprehension and are intended mainly for professional philosophers and serious students. There are also some subsections that most readers can probably scan, or even skip, without significant loss in comprehending the *main* points of the relevant chapter. Technical terms are explained briefly when introduced and are avoided when they can be. Most of the major terms central in epistemology are defined or explicated, and boldfaced numbers in the index indicate main definitional passages. But some are indispensable: they are not mere words, but tools; and some of these terms express concepts valuable outside epistemology and even outside philosophy. The index, by its boldfaced page references to definitions, obviates a glossary.

It should also be stressed that this book is mainly concerned to introduce the *field* of epistemology rather than the *literature* of epistemology—an important but less basic task. It will, however, help non-professional readers prepare for a critical study of that literature, contemporary as well as classical. For that reason, too, some special vocabulary is introduced and a number of the notes refer to contemporary works.

The sequence of topics is designed to introduce the field in a natural progression: from the genesis of justification and knowledge (Part One), to their development and structure (Part Two), and thence to questions about what they are and how far they extend (Part Three). Even apart from its place in this ordering, each chapter addresses a major epistemological topic, and any subset of the chapters can be studied in any order provided some appropriate effort is made to supply the (generally few) essential points for which a later chapter depends on an earlier one.

For the most part this book *does* epistemology rather than talk about it or, especially, about its literature. In keeping with that focus, the ordering of chapters is intended to encourage *understanding* epistemology before

discussing it in large-scale terms, for instance before considering what sort of epistemological theory, say normativist or naturalistic, best accounts for knowledge. My strategy is, in part, to discuss myriad cases of justification and knowledge before approaching analyses of what they are, or the skeptical case against our having them.

In one way, this approach differs markedly from that of many epistemological books. I leave the assessment of skepticism for the last chapter; early passages indicate that skeptical problems must be faced and, in some cases, how they are connected with the subject at hand or are otherwise important. Unlike some philosophers, I do not think extensive discussion of skepticism is the best way to motivate the study of epistemology. Granted, historically skepticism has been a major motivating force; but it is not the only one, and epistemological concepts hold independent interest. Moreover, in assessing skepticism I use many concepts and points developed in earlier chapters; to treat it early in the book, I would have to delay assessing it.

There is also a certain risk in posing skeptical problems at or near the outset: non-professional readers may tend to be distracted, even in discussing conceptual questions concerning, say, what knowledge *is*, by a desire to deal with skeptical arguments purporting to show that there is none. There may be no best or wholly neutral way to treat skepticism, but I believe my approach to it can be adapted to varying degrees of skeptical inclination. An instructor who prefers to begin with skepticism can do so by taking care to explain some of the ideas introduced earlier in the book. The first few sections of Chapter 10 (Chapter 13 in the third edition), largely meant to introduce and motivate skepticism, presuppose far less of the earlier chapters than the later, evaluative discussion; and most of the chapter is understandable on the basis of Part One, which is probably easier reading than Part Two.

My exposition of problems and positions is meant to be as nearly unbiased as I can make it, and where controversial interpretations are unavoidable I try to present them tentatively. In many places, however, I offer my own view. Given the scope of the book, I cannot provide a highly detailed explanation of each major position discussed, or argue at length for my own views. I make no pretense of treating anything conclusively. But in some cases—as with skepticism—I do not want to leave the reader wondering where I stand, or perhaps doubting that there is any solution to the problem at hand. I thus propose some tentative positions for critical discussion.



## Acknowledgments to the first edition

This book has profited from my reading of many articles and books by contemporary philosophers, and from many discussions I have had with them and, of course, with my students. I cannot mention all of these philosophers, and I am sure that my debt to those I will name—as well as to some I do not, such as some whose journal papers I have read but have not picked up again, and some I have heard at conferences—is incalculable. Over many years, I have benefited greatly from discussions with William Alston, as well as from reading his works; and I thank him for detailed critical comments on parts of the manuscript. Reading of books or articles (or both) by Roderick Chisholm, Richard Foley, Paul Moser, Alvin Plantinga, Walter Sinnott-Armstrong, and Ernest Sosa, and a number of discussions with them, have also substantially helped me over many years. My colleagues at the University of Nebraska, especially Albert Casullo, and several of my students have also helped me at many points. I have learned greatly from the participants in the National Endowment for the Humanities seminars and institutes I have directed. I also benefited much from the papers given to the seminars or institutes by (among others) Laurence Bonjour, Fred Dretske, Alvin Goldman, Gilbert Harman, Keith Lehrer, Ruth Marcus, and John Perry, with all of whom I have been fruitfully discussing epistemological topics on one occasion or another for many years.

In relation to some of the main problems treated in the book, I have learned immensely from many other philosophers, including Frederick Adams, Robert Almeder, David Armstrong, John A. Barker, Richard Brandt, Panayot Butchvarov, Carol Caraway, the late Hector-Neri Castañeda, Wayne Davis, Michael DePaul, Susan Feagin, Richard Feldman, Roderick Firth, Richard Fumerton, Carl Ginet, Alan Goldman, Risto Hilpinen, Jaegwon Kim, John King-Farlow, Peter Klein, Hilary Kornblith, Christopher Kulp, Jonathan Kvanvig, Brian McLaughlin, George S. Pappas, John Pollock, Lawrence Powers, W.V. Quine, William Rowe, Bruce Russell, Frederick Schmitt, Thomas Senor, Robert Shope, Donna Summerfield, Marshall Swain, William Throop, Raimo Tuomela, James Van Cleve, Thomas Vinci, Jonathan Vogel, and Nicholas Wolterstorff. In most cases I have not only read some

epistemological work of theirs, but discussed one or another epistemological problem with them in detail.

Other philosophers whose comments or works have helped me with some part of the book include Anthony Brueckner, Stewart Cohen, Earl Conee, Dan Crawford, Jonathan Dancy, Timothy Day, Robert Fogelin, Elizabeth Fricker, Bernard Gert, Heather Gert, David Henderson, Terence Horgan, Dale Jacquette, Eric Kraemer, Noah Lemos, Kevin Possin, Dana Radcliffe, Nicholas Rescher, Stefan Sencerz, James Taylor, Paul Tidman, Mark Timmons, William Tolhurst, Mark Webb, Douglas Weber, Ümit Yalçın, and Patrick Yarnell.

I owe special thanks to the philosophers who generously commented in detail on all or most of some version of the manuscript: John Greco, Louis Pojman, and Matthias Steup. Their numerous remarks led to many improvements. Detailed helpful comments were also provided by readers for the Press, including Nicholas Everett, Frank Jackson, and Noah Lemos. All of the philosophers who commented on an earlier draft not only helped me eliminate errors, but also gave me constructive suggestions and critical remarks that evoked both clarification and other improvements. I am also grateful for permission to reuse much material that appears here in revised form from my *Belief, Justification, and Knowledge* (Wadsworth Publishing Co., 1988) and I thank the editor of *American Philosophical Quarterly* for allowing me to use material from 'The Place of Testimony in the Fabric of Knowledge and Justification' (vol. 34, 1997). For advice and help at several stages I thank Paul Moser, Editor of the series in which this book appears, and Adrian Driscoll and the staff at Routledge in London, including Moira Taylor and Sarah Hall, and Dennis Hodgson.

Robert Audi  
February, 1997

## Preface to the second edition

This preface will presuppose the Preface to the first edition and can therefore be brief. Many improvements have been made in this edition, but they do not make the previous preface inapplicable, and reading it should help anyone considering a study of even part of the book.

My main concern in revising has been to produce a book that is *both* philosophically stronger and easier to read. Doing this has required adding new substantive material, making minor changes throughout, adding or extending many examples, making various refinements and corrections, and bringing in new references, notes, and bibliography.

Instructors who have used the volume in their teaching will find that the content and organization are highly similar and that a transition from the first edition to this one is easy. Students and people reading for general interest should find the book easier to understand. The emphasis is still on enhancing comprehension of the *field* of epistemology—its concepts, problems, and methods—rather than on presenting its literature. But, perhaps even more than in the first edition, the book is generally in close contact with both classical and contemporary literature. In this edition there are also many more references to pertinent books and papers, particularly those published in recent years.

This edition includes more extensive discussion of *virtue epistemology* and *social epistemology*, with *feminist epistemology* figuring significantly (though not exclusively) in relation to social epistemology. The connection of epistemology with philosophy of mind and language also receives more emphasis in this edition. So does *contextualism* and the related theory of “relevant alternatives.”

I am happy to say that Routledge has published a fine and wide-ranging new collection of readings to accompany this book: Michael Huemer’s *Epistemology: Contemporary Readings* (2002). Huemer has chosen classical and contemporary book sections and papers that go well with every chapter in the present book; his larger sections match mine; and he offers helpful introductions to each section and study questions on each chapter.

This edition of my book is certainly self-contained, but its integration with Huemer's supporting collection (for which I have done a long narrative introduction to help both instructors and students) is close, and the two together provide enough substance and diversity to facilitate numerous different kinds of epistemology courses.

## **Acknowledgments to the second edition**

Fortunately, I have continued to benefit from reading of or discussions with most of the people acknowledged above, in the first edition. But since that writing I have come to know the work of many other writers in the field and profited from teaching many more students. I am bound to omit some people I should thank, but I want to acknowledge here a number of people not mentioned above: John Broome, Tyler Burge, David Chalmers, Roger Crisp, Mario DeCaro, Keith DeRose, Rosaria Egidi, Guido Frongia, Douglas Geivett, Joshua Gert, Peter Graham, D.W. Hamlyn, Brad Hooker, Christopher Hookway, Michael Huemer, Jonathan Jacobs, Ralph Kennedy, Simo Knuuttila, Matthias Lutz-Bachmann, Hugh McCann, John McDowell, Tito Megri, Cyrille Michon, Nicholas Nathan, Ilkka Niiniluoto, Tom O'Neil, Derek Parfit, James Pryor, Jlenia Quartarone, Joseph Raz, John Searle, John Skorupski, David Sosa, William Talbot, Wilhelm Vossenkuhl, Fritz Warfield and Paul Weirich. In addition, I have continued to benefit from regular exchanges of ideas or papers (usually both) with many philosophers, including William Alston, Laurence Bonjour, Panayot Butchvarov, Elizabeth Fricker, Alvin Goldman, John Greco, Gilbert Harman, Jaegwon Kim, Christopher Kulp, Jonathan Kvanvig, Bruce Russell, Walter Sinnott-Armstrong, Ernest Sosa, Matthias Steup, Eleonore Stump, Richard Swinburne, Raimo Tuomela, Nicholas Wolterstorff, and Linda Zagzebski.

Readers of earlier versions of this edition deserve special thanks, not only the anonymous readers for the Press but also Michael Pace, Bruce Russell, Mark Owen Webb, and especially Claudio de Almeida, who provided numerous expert comments and criticism (more, indeed, than I could fully take into account in the available time and space). Special thanks are also due the Editors at Routledge, particularly Tony Bruce, Simon Bailey, and Siobhan Pattinson, whose ideas and support have been immensely helpful.

# **Preface to the third edition**

This edition reflects some of the benefits of nearly a decade of teaching and writing about epistemology since the second edition went to press. There are revisions and improvements throughout. Some revisions are responsive to new developments in epistemological literature; others respond to comments by professional colleagues, including various instructors who have used the book. Student responses have also been taken into account.

The book is structurally much as in the second edition, and the previous prefaces largely apply to it. As before, instructors should read prefatory material and the introduction. Long chapters have been divided, but the chapters remain cumulative in content. Most of them, however, can be read independently of the others or by simply looking into some earlier chapter at certain points. The index may also help readers, and its boldface numerals indicate places where the term in question is defined.

Those who have used the second edition will find no difficulty adjusting their teaching or discussions to this one. Chapters that have been divided still cover the same issues, though with new material included and revisions of much that is included. The fit with Michael Huemer's large collection of major papers (cited in the bibliography) is equally good.

Some topics treated in this edition are not addressed in the second. These include the nature of intuitions, the skeptical challenge of rational disagreement, and the value problem: the range of questions concerning why knowledge and justified true belief have value beyond that of merely true belief. Other topics receive considerably more exploration than in the second edition, especially contextualism, perception (including perceptual content), self-evidence and the *a priori*, memorial justification, inferential versus direct knowledge, inference to the best explanation, scientific hypotheses, testimony and trust, understanding, and virtue epistemology.

## Acknowledgments to the third edition

I have been fortunate in being able to exchange ideas with most of the philosophers named in the previous two sets of acknowledgments, and this book has continued to benefit from those discussions. I look back on them with much gratitude. I should particularly mention continuing conversations with the late William P. Alston and with Laurence Bonjour, Mario DeCaro, Elizabeth Fricker, Peter Graham, John Greco, Ralph Kennedy, Peter Klein, Christopher Kulp, Jonathan Kvanvig, Jennifer Lackey, Bruce Russell, Walter Sinnott-Armstrong, Ernest Sosa, and Thomas Vinci.

On one or another topic in the book I have greatly benefited from epistemological discussions with my Notre Dame colleagues Marian David, Michael DePaul, Peter van Inwagen, Alvin Plantinga, Leopold Stubenberg, and Fritz Warfield. Interactions with Andrew Bailey, Michael Bergmann, E.J. Coffman, Roger Crisp, Claudio De Almeida, Keith DeRose, Fred Dretske, Richard Feldman, Richard Fumerton, Alvin Goldman, Christopher Green, Stephen Grimm, Michael Huemer, Thomas Kelly, Matthew Kennedy, Hilary Kornblith, Markus Lamenranta, Duncan Pritchard, Baron Reed, David Sosa, Jeff Speaks, Matthias Steup, Raimo Tuomela, Jonathan Vogel, Michael Zimmerman, and, especially, Sanford Goldberg and Timothy Williamson have also been of great help to me.

Detailed comments on the penultimate version were provided by Scott Hagaman and (on a few parts) by Matthew Lee and Ezra Cook, and these were of great help. Hagaman also made numerous helpful comments and critical remarks in seminar discussions at Notre Dame. The numerous questions and responses by seminar students at the Universities of Notre Dame and Nebraska have been valuable at many points.

The editorial advice of Andrew Beck, Andrew R. Davidson, Michael Andrews, and, since the book's first edition, Tony Bruce have helped me in several important matters. They and their staff have been essential in both design and production.

# Introduction

A sketch of the sources and nature of belief,  
justification, and knowledge

Before me is a grassy green field. A line of trees marks its far edge, which is punctuated by a spruce on its left side and a maple on its right. Birds are singing. A warm breeze brings the smell of roses from a nearby trellis. I reach for a glass of iced tea, still cold to the touch and flavored by fresh mint. I am alert, the air is clear, the scene is quiet. My perceptions are quite distinct.

It is altogether natural to think that from perceptions like these, we come to know a great deal—enough to guide us through much of daily life. But we sometimes make mistakes about what we perceive, just as we sometimes misremember what we have done, or infer false conclusions from what we believe. We may then think we know something when in fact we do not, as when we make errors through inattention or are deceived by vivid dreams. And is it not possible that we are mistaken more often than we think?

## Perception, belief, and justification

Philosophers have thought a great deal about these matters, especially about the nature of perceiving and about what we can know—or may mistakenly think we know—through perception or through other sources of knowledge, such as memory as a storehouse of what we have learned in the past, consciousness as revealing our inner lives, reflection as a way to acquire knowledge of abstract matters, and testimony as providing knowledge originally acquired by others. In approaching these topics in epistemology—the theory of knowledge and justification—it is appropriate to begin with perception. In my opening description, what I detailed was what I *perceived*: what I saw, heard, smelled, felt, and tasted. In describing my experience, I also expressed some of what I *believed*: that there was a green field before me, that there were bird songs, that there was a smell of roses, that my glass felt cold, and that the tea tasted of mint.

It seems altogether natural to believe these things given my experience, and I think I *justifiedly* believed them. I believed them, not in the way I would if I accepted the result of wishful thinking or of merely guessing, but with justification. By that I mean above all that the beliefs I refer to were *justified*.



This a good thing; justified beliefs are of a kind it is desirable and reasonable to hold.

## **Justification as process, as status, and as property**

Being justified, in the sense illustrated by my beliefs about what is clearly before me, need not be the result of a process. Being justified is not, for instance, like being purified, which requires a process of purification. My beliefs about what is before me are not justified because they have been through *a process of being justified*, as when we defend a controversial belief by giving reasons for it. They have not; the question whether they are justified has not even come up. No one has challenged them or even asked why I hold them. They are justified—in the sense that they have the *property of being justified*—*justifiedness*—because there is something about them in virtue of which they are natural and appropriate for me as a normal rational person.

We can see what justifiedness is by starting with a contrast. Unlike believing something that one might arrive at through a wild guess in charades, our justified perceptual beliefs are justified for us simply through their arising in the normal way from our clear perceptions. Roughly, they are justified in the sense that they are quite in order from the point of view of the standards for what we may reasonably believe. That, in turn, is roughly what we may believe without being subject to certain kinds of criticism, say as intellectually lax, as sloppy, as overhasty, or the like. Justified beliefs are also a kind that we tend to expect to be true. Imagine someone's saying 'His belief is justified, but I don't expect it to turn out to be true'. Without special explanation, this would be to take away with one hand something given by the other.

In saying that I justifiedly believe there is a green field before me, I am implying something else, something quite different, though it sounds very similar, namely that I am *justified in believing* there is a green field before me. To see the difference, notice that we can be justified *in* believing something—roughly in the sense that we *have* a justification for believing it—without believing it at all, quite as we can be justified in doing something, such as criticizing a person who has failed us, without doing it. Similarly, I might be justified in believing that I can do a certain difficult task, yet fail to believe this until someone helps me overcome my hesitation. I may then see that I *should* have believed it.

Being justified in believing something is having justification *for* believing it. This, in turn, is roughly a matter of having ground for believing it (and we also speak of having *a* ground or *a* justification or *a* reason). Just as we can have reason to do things we do not do, we can have reason to believe things we do not believe. You can have reason to go to the library and forget to, and I can have reason to believe someone is making excuses for me but—because I have no inkling that I need any—fail to believe this. Our justification for believing is basic raw material for actual justified belief; and justified belief is commonly good raw material for knowledge.

The two justificational notions are intimately related: if one justifiedly believes something, one is also justified *in* believing it, hence *has* justification for believing it. But the converse does not hold: not everything we are justified *in* believing is something we *do* believe. When I look at a lawn, I am justified *in* believing it has more than ten blades of grass per square foot, but I would not normally have any belief about the number of blades per square foot. We have more justificational raw material than we need or use. We do not believe anywhere near the number of things that we have justification to believe. This holds not just in trivial matters but also in, for instance, mathematics.

There are many things we are justified in believing which we do not actually believe, such as the proposition that normal people do not drink 100 liters of water a day. Let us call the first kind of justification—justifiedly believing—*belief justification*, as it belongs to actual beliefs. It is also called *doxastic* justification, from the Greek *doxa*, translatable as ‘belief’. Call the second kind—being justified *in* believing—*situational justification*, since it is based on the informational situation one is in. It is a *status* one has in virtue of that situation. This situation includes not just what one perceives, but also one’s background beliefs and knowledge, such as the belief that people drink at most a few liters of water a day. Situational justification is also called *propositional justification*, since the proposition in question is justified *for* the person whose situation provides justification for believing it, and the person *has* justification for it.

In any ordinary situation in waking life, we have both a lot of general information stored in memory and much specific information presented in our perceptions. We do not need all this information, and our situational justification for believing something is often unaccompanied by our actually believing that it is so. We have situational justification for vastly more justified beliefs than we actually have. Here nature is very generous. We are built to gain from a mere glance enough information to ground vastly more beliefs than we normally form or rely on.

Without situational justification, such as the kind that comes from seeing a green field, there would be no belief justification. I would not, for instance, justifiedly believe that there is a green field before me. We cannot have a justified belief without being in a *position* to have it. Without situational justification, we are not in such a position. Without belief justification, on the other hand (i.e., doxastic justification), we would have no beliefs of a kind we want and need, those with a positive status—being justified—that makes them appropriate for us as rational creatures and warrants us in expecting them to be true. Belief justification, then, is more than the situational kind it presupposes.

Belief justification occurs when there is a certain kind of connection between what yields situational justification and the justified belief that benefits from it. Belief justification occurs when a belief is *grounded in*, and thus in a way supported by (or *based on*), something that gives one situational

justification for that belief, such as seeing a field of green. Seeing is of course perceiving; and perceiving is a basic source of knowledge—perhaps our most elemental source, at least in childhood. This is largely why perception is so large a topic in epistemology.

## Knowledge and justification

Knowledge would not be possible without belief justification—or a kind of grounding significantly like it. If I did not have the kind of justified belief I do—if, for instance, I were wearing dark sunglasses and could not tell the difference between a green field and a smoothly ploughed one that is really an earthen brown—then on the basis of what I now see I would not know that there is a green field before me.

To see how knowledge fits into the picture so far sketched, consider two points. First, justified belief is important for knowledge because at least the typical things we know we *also* justifiedly believe on the same basis that grounds our knowing them. If I know someone is making excuses for me, say by the way she explains my lateness, I do not just believe this but justifiedly believe it. Second, much of what we justifiedly believe we also know. Surely I could have maintained, regarding each of the things I have said I justifiedly believed through perception, that I also knew it. And do I not know these things—say that there is a lawn before me and a car on the road beyond it—on the *same* basis on which I justifiedly believe them, for instance on the basis of what I see and hear? This is very plausible.

As closely associated as knowledge and justified belief are, there is a major difference. If I know that something is so, then it is *true*, whereas I can justifiedly believe something false. If a normally reliable friend tricked me into believing something false, say that he lost my car keys, I could still justifiedly believe he lost them. We must not assume, then, that everything we learn about justified belief applies to knowledge. We should look at both concepts independently.

I said that I *saw* the green field and that my belief that there was a green field before me arose from my seeing it. If the belief arose, under normal conditions, from my seeing the field (so that I believed it is there simply because I saw it there), then the belief was true, justified, and constituted knowledge. Again, however, we can alter the example to bring out how knowledge and justification may diverge: the belief might remain justified even if, unbeknownst to me, the grass had been burned up since I last saw it, and there was now a perfect artificial replica of it spread out in grassy-looking strips of cloth that hide the charred ground. Then, although I might think I know the green field is there, I would only falsely believe I know this. Such a bizarre happening is, to be sure, improbable. Still, a justified but false belief *could* arise in this way.

## Memory, introspection, and self-consciousness

As I look at the field before me, I *remember* carefully cutting a poison ivy vine from the trunk of the spruce. Surely, my memory belief that I cut off this vine is justified. I think I also know that I did this. But here I confess to being less confident than I am of the justification of my perceptual belief, held in the radiant sunlight, that there is (now) a green field before me.

As our memories become less vivid, we tend to be correspondingly less sure that our beliefs apparently based on them are justified. Still, I distinctly recall cutting the vine. The stem was furry; it was bonded to the tree trunk; the cutting was difficult and slightly wounded the tree. By contrast, I have no belief about whether I did this in the summer or the fall. I *entertain* the proposition that it was in the summer; I *consider* whether it is true; but, being utterly uncertain, I *suspend judgment* on it. I thus neither believe it nor *disbelieve* it, that is, believe it is false. My stance is one of *non-belief*. I need not try to force myself to resolve the question and judge the proposition either way. I might need to resolve it if something important turned on when I did the pruning; but here suspended judgment, with the resulting non-belief, is not uncomfortable.

As I think about cutting the vine, it occurs to me that in recalling that task, I am vividly imaging it. Here, I seem to be looking into my own consciousness, thus engaging in a kind of *introspection*. I can still see, in my mind's eye, the furry vine clinging to the tree, the ax, the sappy wound along the trunk where the vine was severed from it. I have turned my attention inward to my own imagery. The object of my attention, my own imaging of the scene, seems internal and is present to my consciousness, though *its* object is external and long gone by. But clearly, I believe that I am imaging the vine; and there is no apparent reason to doubt that I justifiedly believe this and know that it is so. This is a simple case of *self-knowledge*.

## Reason and rational reflection

I now look back at the field and am struck by how perfectly rectangular it looks. If it is perfectly rectangular, then its corners are right angles. Here I believe something different in kind from the things cited so far: that if the field is rectangular, then its corners are right angles. This is a geometrical belief. I do not hold it on the same sort of basis I have for the other things I have mentioned believing. My conception of geometry as applied to ideal figures seems to be my basis. On that basis, my belief seems to be firmly justified and to constitute knowledge.

I can see that the spruce is taller than the maple, and that the maple is taller than the crab apple tree on the lawn closer by. I now realize that the spruce is taller than the crab apple. My underlying belief here is that if one thing is taller than a second and the second taller than a third, then the first is taller than the third. And, perhaps even more than the geometrical belief,

this abstract belief seems to arise simply from my grasp of the concepts in question, above all the concept of one thing's being taller than another.

## Testimony

The season has been dry, and it now occurs to me that the roses will not flourish without a good deal of water. But this I do not believe simply on the basis of perception. One source from which I learned it is repeated *observation*. But there is another possible source: although much knowledge comes directly from our own experience, much also originates with *testimony* from others. I have received testimony as to where on the stem to trim off dead roses. If I did not learn about watering roses from my own experience, I could have learned the same things from testimony, just as I learned from a friend how far back to clip off dead roses.

To be sure, I need perception, such as hearing what I am told, to acquire knowledge on the basis of testimony, just as I needed perception to *learn* these things about roses on my own; and I need memory to *retain* them whatever their source. They are, however, generalizations and hence do not arise from perception in the direct and apparently simple way my visual beliefs do, or emerge from memory in the way my beliefs about past events I witnessed do. But do I not still justifiedly believe that the roses will not flourish without a lot of water? The commonsense view is that I both justifiedly believe and know this about roses, and that I can know it either through generalizing—a kind of reasoning—from my own observations, or from testimony, or from both.

## Basic sources of belief, justification, and knowledge

The examples just given represent what philosophers have called perceptual, memorial, introspective, a priori, inductive, and testimony-based beliefs. The first four kinds are basic in epistemology. My belief that the glass is cold to the touch is *perceptual*, being based as it is on tactual perception. My belief that I cut the poison ivy vine from the spruce is *memorial*, since it is stored in my memory and held because of that fact. My belief that I am imagining a green field is called *introspective* because it is conceived as based on “looking within” (the etymological meaning of ‘introspection’); but it could also be called simply *self-directed*: no “peering” within or special concentration is required. My belief that if the spruce is taller than the maple and the maple is taller than the crab apple then the spruce is taller than the crab apple is called *a priori* (meaning, roughly, based on what is “prior” to observational experience) because it apparently arises not from experience of how things actually behave, but simply in an *intuitive* way. It arises from a rational grasp of the key concepts one needs in order to have the belief, such as the concept of one thing's being taller than another.

By contrast, my belief that the roses will not grow well without abundant

water does not arise directly from one of the four basic sources just mentioned: perception, memory, introspection, and a priori intuition (reason, in one sense of the term). It is called *inductive* because it is formed (and held) on the basis of a generalization from something more basic, in this case what I learned from perceptual experiences with roses. Those experiences, apparently through my beliefs recording them, “lead into” the generalization about roses, to follow the etymological meaning of ‘induction’. For instance, I remember numerous cases in which roses have faded when dry, and I eventually concluded that they need abundant water.

Each of the four basic kinds of belief I have described—perceptual, memorial, introspective, and a priori—is grounded in the source from which it arises. The nature of this grounding is explored in detail in Part One, which concerns perception, memory, consciousness, and reason. These sources are commonly taken to provide raw materials for inductive generalizations, as where observations and memories about roses yield a basis for generalizing about their needs.

Any of the beliefs we considered could instead have been grounded in testimony (the topic of Chapter 7), had I formed the beliefs on the basis of being given the same information by someone I trust. That person, however, would presumably have acquired it through one of these *other* sources (or ultimately through someone’s having done so), and this makes testimony a different kind of source. This is why testimony is not a basic source of knowledge. It is still, however, incalculably important for human knowledge and unlimitedly broad. It can, for instance, justify a much wider range of propositions than perception can. We can credibly tell others virtually anything we know.

### Three kinds of grounds of belief

Our examples illustrate not only grounding of beliefs *in* a source, such as perception or introspection, but also *how* they are grounded in these sources. There are at least three important kinds of grounding of beliefs—ways they are grounded. These are causal, justificational, and epistemic grounding. All three are important for many major epistemological questions.

Consider my belief that there is a green field before me. It is *causally grounded* in my experience of seeing the field because that experience produces or underlies the belief. It is *justificationally grounded* in that experience because the experience, or at least some element in the experience, justifies my belief. And it is *epistemically grounded* in the experience because in virtue of that experience my belief constitutes knowledge that there is a green field before me (‘epistemic’ comes from the Greek *episteme* meaning, roughly, ‘knowledge’). These three kinds of grounding very often coincide (though Chapter 11 will describe important cases in which knowledge and justification do not). I will thus often speak simply of a belief as *grounded* in a source, such as visual experience, when what grounds the belief does so in all three ways.

Causal, justificational, and epistemic grounding each go with a very common kind of question about belief. Let me illustrate.

Causal grounding goes with ‘Why do you believe that?’ An answer to this, asked about my belief that there is a green field before me, would be that I see it. This is the normal kind of reply; but as far as mere causal production of beliefs goes, the answer could be brain manipulation or mere hypnotic suggestion. If, however, mere brain manipulation or mere hypnotic suggestion produces a belief, then the causal ground of the belief would not justify it. If, under hypnosis, I am told that someone dislikes me and as a result I believe this, the belief is not thereby justified.

Justificational grounding goes with such questions as ‘What is your justification for believing that?’ or ‘What justifies you in thinking that?’ or ‘Why should I accept that?’ (‘Why do you believe that?’ *can* be asked with this same justification-seeking force.) Again, I might answer that I see it. I might, however, have a justification (the situational kind) that, unlike seeing the truth in question, is not a *cause* of my believing it.

The justification I cite could also be the testimony of a credible good friend. It could be this even when, by a short circuit, brain manipulation does the causal work of *producing* my belief and leaves the testimony like a board that slides just beneath a roof beam but bears none of its weight. This shows that an element that provides situational justification for a belief may play no role in producing or supporting the belief, even if this element, like the auxiliary unstressed board, stands ready to play a supporting role if the belief is put under pressure by a challenge.

Epistemic grounding goes with ‘How do you know that?’ Once again, saying that I see it will commonly answer this. Here, however, it may be that a correct answer *must* cite something that is *also* a causal ground for the belief (a matter discussed in Chapter 10). Certainly a justificational ground need not be a ground of knowledge. One can justifiably believe a proposition without knowing it.

Clearly, the same sorts of points can be made for the other five cases I have described: memorial beliefs are grounded in memory, self-directed (“introspective”) beliefs in consciousness, inductively based beliefs in further, premise-beliefs that rest on experience, a priori beliefs in reason, and testimony-based beliefs in testimony.

## Fallibility and skepticism

Even well-grounded beliefs can be mistaken. We can be deceived by our senses. We are fallible in perceptual matters, as in our memories, in our reasoning, and in other respects. One might now wonder, as skeptics do, whether we *know* even that it is improbable that our senses are now deceiving us. One might also wonder whether, when we take ourselves to see green grass, we are even *justified* in our belief that no such mistake has occurred.

Suppose that I am in an unfamiliar park. I might not know or even



justifiedly believe that artificial grass has not replaced the natural grass I take to be before me. (I may have heard of such substitutions and may have no good reason to believe this has not happened, though I do not consider the matter.) Am I justified in believing that there is green grass before me?

Suppose that I am not justified in believing there is green grass before me. If not, how can I be justified in believing what appear to be far less obvious truths, such as that my home is secure against the elements, my car safe to drive, and my food free of poison? And how can I know the many things I need to know in life, such as that my family and friends are trustworthy, that I can control my behavior and thus partly determine my future, and that the world we live in at least approximates the structured reality portrayed by common sense and science?

These are difficult and important questions. They indicate how insecure and disordered human life would be if we could not suppose that we possess justified beliefs and knowledge. We stake our lives every day on what we take ourselves to know. It would be unsettling to revise this stance and retreat to the view that at best we have justification to believe. But if we had to give up even this moderate view and to conclude, say, that what we believe is not even justified, we would face a crisis. Much later, in discussing skepticism, I will explore such questions at some length. Until then I will assume the commonsense view that beliefs with a basis like that of my belief that there is a green field before me are not only justified but also constitute knowledge.

Once we proceed on this commonsense assumption, it is easy to see that there are many different kinds of circumstances in which beliefs arise in such a way that they are apparently both justified and constitute knowledge. In considering this variety of circumstances yielding justification and knowledge, we can explore how beliefs are related to perception, memory, consciousness, reason, and testimony (the topics of Chapters 1–7).

## Overview

There is a great deal more to be said about each of these sources of belief, justification, and knowledge and about how they ground what they do ground. The first seven chapters explore, and in some cases compare, the basic sources of belief, justification, and knowledge.

In the light of what those chapters show, we can discuss the development and structure of knowledge and justification (the task of Part Two). Much of what we believe does not come directly from perception, memory, introspection, or reflection of the kind appropriate to knowledge of such truths as those of elementary mathematics or those turning on our grasp of simple relations, for instance the proposition that if the spruce is taller than the maple, then the maple is shorter than the spruce, which we know by virtue of understanding the relations expressed by ‘taller’ and ‘shorter’. We must explore how inference and other developmental processes expand our body of knowledge and justified beliefs (this is the task of Chapter 8). Moreover,



once we think of a person as having the resulting complex body of knowledge and justified belief, we encounter the questions of what structure that large and intricate body has, and of how its structure is related to the amount and kind of knowledge and justification it contains. As we shall see in Part Two, these structural questions take us into an area where epistemology and the philosophy of mind often overlap.

On the basis of what Part One shows about sources of knowledge and justification and what Part Two shows about their development and structure, we can fruitfully proceed to consider more explicitly what knowledge and justification are and what kinds of things can be known (the task of Part Three). It is true that if we had no sense at all of what they are, we could not find the kinds of examples of them needed to explore their sources and their development and structure. If we do not have before us a wide range of examples of justification and knowledge, we lack the data appropriate to seeking a philosophically illuminating analysis of them. It is in the light of the examples and conclusions of Parts One and Two that Chapters 10 and 11 clarify the concept of knowledge, and, to a lesser extent, that of justification, in some detail.

With a conception of knowledge laid out, it is possible to explore the apparent extent of knowledge and justification in three major territories—the scientific, the ethical, and the religious. In exploring these domains, Chapter 12 applies some of the epistemological results of the earlier chapters. These chapters continue to assume the commonsense view that we have a great deal of knowledge and justification. If, however, skepticism is justified, then the commonsense assessment that the first twelve chapters make regarding the extent of knowledge and justification must be revised. Whether skepticism is justified is the focus of Chapters 13 and 14.

Along the way in all fourteen chapters, there is much to be learned about concepts that are important both in and outside epistemology, especially those of belief, causation, certainty, coherence, explanation, fallibility, illusion, inference, intellectual virtue, introspection, intuition, meaning, memory, rationality, reasoning, relativity, reliability, truth, and understanding. There are also numerous epistemological positions to be considered, sometimes in connection with historically influential philosophers. But the main focus will be on the major concepts and problems in the field, not on any particular philosopher or text. This may well be the best way to facilitate studying philosophers and epistemological texts; it will certainly simplify an already complex task.

Knowledge and justification are not only interesting in their own right as central epistemological topics; they also represent positive values in the life of every reasonable person. For all of us, there is much we want to know. We also care whether we are justified in what we believe—and whether others are justified in what they tell us. The study of epistemology can help in making this quest, even if it often does so indirectly. It can certainly help us assess how well we have done in the quest when we review our results.

Well-developed concepts of knowledge and justification can serve as ideals in human life. Positively, we can try to achieve knowledge and justification in relation to subjects that concern us. Negatively, we can refrain from forming beliefs where we think we lack justification, and we can avoid claiming knowledge where we think we can at best hypothesize. If we learn enough about knowledge and justification conceived philosophically, we can better search for them in matters that concern us and can better avoid the dangerous pitfalls that come from confusing mere impressions with justification or mere opinion with knowledge. This is not to say that epistemological knowledge can be guaranteed to yield new everyday knowledge. But the more we know about the constitution of knowledge and justification, the better we can build them through our own inquiries, and the less easily we will fall into the pervasive temptation to take an imitation to be the real thing.



## **Part One**

Sources of justification,  
knowledge, and truth



# I Perception

Sensing, believing, and knowing

- **The elements and basic kinds of perception**

*Perceptual belief*

*Perception, conception, and belief*

*Propositional and objectual perception*

- **Seeing and believing**

*Perceptually embedded beliefs*

*Perception as a source of potential beliefs*

*The perceptual hierarchy*

*Simple, objectual, and propositional perception*

*The informational character of perception*

- **Perceptual justification and perceptual knowledge**

*Seeing and seeing as*

*Perceptual content*

*Seeing as and perceptual grounds of justification*

*Seeing as a ground of perceptual knowledge*

# I Perception

## Sensing, believing, and knowing

As I look at the green field before me, I might believe not only that there is a green field there but also that I *see* one. And I do see one. I visually perceive it. Both beliefs, the belief that there is a green field there, and the self-referential belief that *I* see one, are grounded, causally, justificationally, and epistemically, in my perceptual experience. They are produced by that experience, justified by it, and constitute knowledge in virtue of it.

The same sort of thing holds for the other senses. Consider touch. I not only believe, through touch (as well as sight), that there is a glass here, I also feel its cold surface. Both beliefs—that there is a glass here and that it is cold—are grounded in my tactual experience. I could believe these things on the basis of someone's testimony. My beliefs would then have a quite different status. For instance, my belief that there is a glass here would not be a *perceptual belief*, but only a belief *about a perceptible*, that is, a perceivable object, the kind of thing that can be seen, touched, heard, smelled, or tasted. Through testimony we have beliefs about perceptibles we have never seen or experienced in any way.

My concern is not with the hodgepodge of beliefs that are simply about perceptibles, but with perception and perceptual beliefs. Perceptual beliefs are not simply beliefs about perceptibles; they are *beliefs grounded in perception*. We classify beliefs as perceptual by the nature of their roots, not by the color of their foliage; by their grounds, not their type of content. Those roots may be visual, auditory, and so forth for each perceptual mode. But vision and visual beliefs are an excellent basis for discussing perception, and I will concentrate on them and mention the other senses only when it adds clarity.

Perception is a source of knowledge and justification mainly by virtue of yielding beliefs that *constitute* knowledge or are justified. But we cannot hope to understand perceptual knowledge and justification simply by exploring those beliefs. We must also understand what perception is and *how* it yields beliefs. We can then begin to understand how it yields knowledge and justification or—sometimes—fails to yield them.

## The elements and basic kinds of perception

There are apparently at least four elements in perception: (1) the perceiver, me; (2) the object, the field I see; (3) the sensory experience, say my visual experience of colors and shapes; and (4) the relation between the object and the subject, commonly taken to be a causal relation by which the object produces the sensory experience in the perceiver. To see the field is apparently to have a certain sensory experience as a result of the impact of the field on our vision.

Some accounts of perception add to the four items on this list; others subtract from it. To understand perception we must consider both kinds of account and how these elements are to be conceived in relation to one another. But first, it is essential to explore examples of perception.

There are several quite different ways to speak of perception. Each corresponds to a different way of perceptually responding to experience. We often speak simply of what people perceive, for instance see. We also speak of what they perceive the object to be, and we commonly talk of facts they know through perception, such as that the grass is long. Visual perception most readily illustrates this, so let us start there.

I see, hence perceive, the green field. Second, speaking in a less familiar way, I see it *to be* rectangular. Thus, I might say that I know it looks irregular from the nearby hill, but from the air you can see it to be perfectly rectangular. Third, I see *that it is* rectangular. Perception is common to all three cases. Seeing, which is a paradigm perception, is central in each.

The first case is one of *simple perception*, perception taken by itself (here, visual perception). I see the field, and this experience is the visual parallel of hearing a bird (an auditory experience), touching a glass (a tactual experience), smelling roses (an olfactory experience), and tasting mint (a gustatory experience). If the first case is simply *perceiving of* some object, the second is a case of *perceiving to be*, as it is seeing something to be so: I do not just see the field, as when I drive by at high speed and do not even realize what is in my peripheral vision; rather, I see the field to be rectangular. The third case is one of *perceiving that*; it is seeing that a particular thing is so, namely that the field is rectangular.

These cases represent three kinds, or *modes*, of perception. Perception of the simplest kind (or in the simplest mode), such as seeing, occurs in all three; but, especially because of their relation to knowledge and justified belief, they are significantly different. We can best understand these three kinds (or modes) of perception if we first focus on their relation to belief.

### *Perceptual belief*

The last two cases—perceiving *that*, and perceiving *to be*—are different from the first—perceiving *of*—in implying corresponding kinds of beliefs: seeing that the field is rectangular implies believing that it is, and seeing it



to be green implies believing it to be green. If we consider how both kinds of beliefs—beliefs *that* something is so and beliefs *of* (hence *about*) something—are related to perception, we can begin to understand how perception occurs in all three cases, the simple and the more complex. In my second and third examples of perception, visual perception (seeing) issues in beliefs that are grounded in seeing and can thereby constitute visual knowledge, such as knowing that the field is green.<sup>1</sup>

In our example of simple perception, my just seeing the field provides a basis for both kinds of beliefs. It does this even if, because my mind is entirely occupied with what I am hearing on the radio as I glance over the field, no belief about the field actually arises in me. The visual experience is, in this instance, like a foundation that has nothing built on it but is ready to support a structure. If, for example, someone were to ask if the field has shrubbery, then given the lilacs prominent in one place, I might immediately form the belief that it does and assent. This belief is visually grounded; it comes *from* my seeing the field though it did not initially come *with* it. When visual experiences do produce beliefs, as they usually do, what kinds of beliefs are these, and how are they specifically perceptual?

Many of my beliefs arising through perception correspond to perception *that*, say to seeing that the lilacs are blooming. I believe that the field is lighter green toward its borders, that it is rectangular in shape, and that it has many ruts. But I may also have various beliefs about it that are of the second kind: they correspond to perception *to be*, for instance to seeing something to be a certain color. Thus, I believe the field to be green, to be rectangular, and so on. The difference between these two kinds of belief is significant. As we shall shortly see, it corresponds first of all to two distinct ways in which we are related to the objects we perceive and, second, to two different ways of assessing the truth of what, on the basis of our perceptions, we believe.

The first kind of belief just described is the kind people usually think of when they consider beliefs: it is called *propositional*, as it is generally considered a case of believing a proposition—say, *that* the field is rectangular. The belief is thus true or false depending on whether the proposition in question—here *that the field is rectangular*—is true or false. In holding the belief, moreover, in some way I think of what I see as a field which is rectangular: in believing that the field is rectangular, I *conceive* what I take to be rectangular as a field.

The second kind of belief might be called *objectual*: it is a belief regarding an object, say the field, with which the belief is actually connected. This is an object *of* (or about) which I believe something, say that it is rectangular. If I believe the field to be rectangular, there really is such an object, and I have a certain relation to it. A special feature of this relation is that there is no particular proposition I must believe about the field. To see that there is no particular proposition, notice that in holding this objectual belief I need not think of what I see *as* a field. I might mistakenly take it to be (for instance) a lawn or a grasslike artificial turf, yet still believe it to be rectangular. I might

think of it just in terms of what I believe it to be and not in terms of anything else.

Thus, although there is *some* property I must take the field to have—corresponding to what I believe it to *be*—there is no other particular way I must think of it. With objectual belief, then, there is no particular notion, no specific conceptual “handle,” that must yield the subject of any proposition I believe about the object: I do not have to believe that the *field* is green, that the *grass* is green, or any such thing. Perception leaves us vast latitude as to what we learn from it. People differ greatly in the beliefs they form about the very same things they see.<sup>2</sup>

The concept of objectual perception, then, is very permissive about what one believes about the object perceived. This is one reason why it leaves so much space for imagination and learning—a space often filled by the formation of propositional beliefs, each capturing a different aspect of what is perceived, say that the field is richly green, that it is windblown, and that it ends at a treeline.

A different example may bring these points out further. After seeing a distant flare and coming to believe, of something blurry and far away, that it glowed, one might ask, ‘What on Earth was it that glowed?’ Before we can believe the proposition that a *flare* glowed, we may have to think about where we are, the movement and fading of the glow, and so forth. The objectual belief is a guide by which we may arrive at propositional beliefs and propositional knowledge.

### *Perception, conception, and belief*

The same kind of example can be used to illustrate how belief depends on our conceptual resources in a way that perception does not. Suppose I had grown up in the desert and somehow failed to acquire the concept of a field. I could nonetheless see the green field, and from a purely visual point of view it might *look* the same to me as it does now. I could also believe, regarding the field I see—and perhaps conceive as sand artificially covered with something green—that it is rectangular. But I could not believe that the *field* is rectangular. This propositional belief as it were portrays what I see *as* a field in a way that requires my having a concept of one.

There is a connection here between thought and language (or at least conceptualization). If I believe (think) that the field is rectangular, or even simply have the thought that it is, I should be able to *say* that it is and to know what I am talking about. But if I had no concept of a field, then in saying this I would not know what I am talking about.<sup>3</sup> Similarly, a two year old, say, Susie, who has no notion of a tachistoscope, can, upon seeing one and hearing its fan, believe it to be making noise; but she cannot believe specifically that the tachistoscope is making noise. Her propositional belief, if any, would be, say, that the thing on the table is making noise. Since this is true, what she believes is true and she may know this truth, but she need not know much

about the object this truth concerns: in a way, she does not know what it is she has this true belief *about*.

The general lesson here is important. A basic mode of learning about objects is to find out truths about them in this elementary way: we get a handle on them through perceptually discriminating some of their properties; we form objectual (and other) beliefs about them from different perspectives; and (often) we finally reach an adequate concept of what they are. From the properties I believe the flare in the distance to have, I finally figure out that it is a flare that has them. This suggests that there is at least one respect in which our knowledge of (perceptible) *properties* is more basic than our knowledge of the substances that have them; but whether that is so is a question I cannot pursue here.

Unlike propositional beliefs, objectual beliefs have a significant degree of indefiniteness in virtue of which it can be misleading simply to call them true or false; they are accurate or inaccurate, depending on whether what one believes of the object (such as that it is rectangular) is or is not *true of it*. Recall Susie. If she attributes noise-making to the tachistoscope, she truly believes, *of it*, that it is making noise. She is, then, *right about it*. But this holds even if she has no specific concept of what it is that is making the noise. If we say unqualifiedly that her belief about it is true, we invite the question 'What belief?' and the expectation that the answer will specify a particular proposition, say that the tachistoscope is making noise. But it need not, and we might be unable to find any proposition that she does believe about it. She can be right about something without knowing or even having any conception of what kind of thing it is that she is right about.

Knowledge is often partial in this way. Still, once we get the kind of epistemic handle on something that objectual belief can provide, we can usually use that to learn more about it.<sup>4</sup> Suppose I see a dog's tail projecting from under a bed and do not recognize it as such. If I believe it to be a slender furry thing, I have a place to start in finding out what else it is. I will, moreover, be disposed to form such beliefs as that there is a slender furry thing there. I will also have justification for them. But I need not form them, particularly if my attention quickly turns elsewhere.

### *Propositional and objectual perception*

Corresponding to the two kinds of beliefs I have described are two ways of talking about perception. I see *that* the field is rectangular. This is (visual) *propositional perception*: perceiving *that*. I also see it *to be* rectangular. This is (visual) *objectual perception*: perceiving *to be*. The same distinction apparently applies to hearing and touch. Perhaps, for example, I can hear that a piano is out of tune by hearing its sour notes, as opposed to hearing the tuner say it needs tuning. As for taste and smell, we speak as if they yielded only simple perception: we talk of smelling mint in the iced tea, but not of smelling that it is minty or smelling it to be minty. Such talk is, however,

intelligible on the model of seeing that something is so and seeing it to be so. We may thus take the distinction between perceiving *that* and perceiving *to be* to apply in principle to all the senses.

It is useful to think of perceptual beliefs as *embedded* in the corresponding propositional or objectual perception, roughly in the sense that they are integrally tied to perceiving of that kind and derive their character and perhaps their authority from their perceptual grounding. Take propositional belief first. My belief that the field is rectangular is embedded in my seeing that it is, and Susie's believing the tachistoscope to be making noise is embedded in her hearing it to be doing so. In each case, the belief is an element in perception of the corresponding kind. These kinds of perception might therefore be called *cognitive*, since belief is a cognitive attitude: roughly the kind having a proposition (something true or false) as its object.<sup>5</sup> The object of the belief that the field is rectangular is the specific proposition that the field is rectangular, which is true or false.

Now consider objectual perceptual beliefs. If believing the tachistoscope *to be* making noise has a propositional object, that object may be plausibly taken to be some proposition or other to the effect that it is making noise, which (though left unspecified by the ascription of the belief) is also true or false. But some objectual perceptions may also be plausibly conceived as simply attributions of a perceptible property to the thing perceived; here the embedded objectual belief is *true of* the object rather than simply true. A tiny, prelingual child might see the liquid offered to it to be milk yet not believe (or disbelieve) the proposition that it is milk. In this respect, belief is unlike attitudes of approval or admiration or indignation, which are evaluated not as true or false but rather as, say, appropriate or inappropriate.<sup>6</sup>

Both propositional and objectual beliefs are grounded in simple perception. If I do not see a thing at all, I do not see that it has any particular property and I do not see it to be anything. Depending on whether perceptual beliefs are propositional or objectual, they may differ in the kind of knowledge they give us. Propositional perception yields knowledge both of *what* it is that we perceive and of some *property* of it, for instance of the *field's* being *rectangular*. Objectual perception may, in special cases, give us knowledge only of a property of what we perceive, say of its being green, when we do not know what it is or have any belief as to what it is.

In objectual perception, we are, to be sure, in a good position to come to know *something* or other about the object, say that it is a green expanse. Objectual perception may thus give us information not only about objects of which we have a definite conception, such as home furnishings, but also about utterly unfamiliar objects of which we have at most a very general conception, say 'that noisy thing'. This is important. We could not learn as readily from perception if it gave us information only about objects we conceive in the specific ways in which we conceive most of the familiar things we see, hear, touch, taste, and smell.<sup>7</sup>

## Seeing and believing

Both propositional and objectual perceptual beliefs are commonly grounded in perception in a way that apparently connects us with the outside world and assures their truth. For instance, my visual belief that the field is rectangular is so grounded in my seeing the field that I veridically (truly) see that it is rectangular; my tactually believing the glass to be cold is so grounded in my feeling it that I veridically feel it to be cold. Let us explore the relation between perception and belief.

### *Perceptually embedded beliefs*

Must beliefs grounded in seeing be true? Admittedly, I might visually (or tactually) believe that something is rectangular under conditions poor for judging it. Compare viewing a straight stick half submerged in water (it will look bent). My visually grounded belief might then be mistaken. But such a mistaken belief is not *embedded* in propositional perception that the stick is bent—that proposition is false and hence is *not* something one sees is so (or to be so). The belief is merely *produced* by some element in the simple perception of the stick: I see the stick in the water, and the operation of reflected light causes me to have the illusion of a bent stick. I thus do not see that the stick is bent: my genuine perception is of it, but not of its curvature. Seeing that curvature or seeing that the stick is bent would entail that it *is* bent, which is false. If the stick is not bent, I cannot see that it is.

As this suggests, there is something special about both perceiving *that* and perceiving *to be*. They are *veridical experiences*, that is, they imply truth. Specifically, if I see that the field is rectangular, or even just see it to be rectangular, then it truly is rectangular. Thus, when I simply see the rectangularity of the field, if I acquire the corresponding embedded perceptual beliefs—if I believe that it is rectangular when I see that it is, or believe it to be rectangular when I see it to be—then I am correct in so believing.

Perceiving *that* and perceiving *to be*, then, imply (truly) believing something about the object perceived—and so are *factive*. Does simple perception, perception *of* something, which is required for either of these more complex kinds of perception, also imply true belief? Very commonly, simple perception does imply truly believing something about the object perceived. If I hear a car go by, I commonly believe a car is passing. But could I not hear it, but be so occupied with my reading that I form no belief about it? Let us explore this.

### *Perception as a source of potential beliefs*

As is suggested by the case of perception overshadowed by preoccupation with reading, there is reason to doubt that simple perceiving *must* produce any belief at all. Moreover, it commonly does not produce beliefs even of

what *would* be readily believed if the question arose. Suppose I am looking appreciatively at a beautiful rug. Must I believe that it is not producing yellow smoke, plain though this fact is? I think not; there seems to be a natural economy of nature—perhaps explainable on an evolutionary basis—that prevents our minds from being cluttered with the innumerable beliefs we would have if we formed one for each fact we can see to be the case.

This line of thought may seem to fly in the face of the adage that seeing is believing. But properly understood, that may apply just to propositional or objectual seeing. In those cases, perception plainly does entail beliefs. Seeing that golf ball-size hail is falling *is* (in the sense that it entails) believing it.<sup>8</sup> This fact, however, is not only perceptible; it is striking.

In any event, could I see the field and believe nothing regarding it? Must I not see it to be something or other, say green? And if so, would I not believe, of it, *something* that is true of it, even if only that it is a green object some distance away? Consider a different example.

Imagine that we are talking excitedly and a bird flies quickly across my path. Could I see it, yet form no beliefs about it? There may be no clearly correct answer. For one thing, although there is much we *can* confidently say about seeing and believing, ‘seeing’ and ‘believing’ are, like most philosophically interesting terms, not precise. They have an element of vagueness. No standard dictionary definition or authoritative statement can be expected either to tell us precisely what they mean or, especially, to settle every question about when they do and do not apply.<sup>9</sup> Still, we should be wary of concluding that vagueness makes any significant philosophical question unanswerable. How, then, should we answer the question whether seeing entails believing?

A negative response might be supported as follows. Suppose I merely see the bird but pay no *attention* to it because I am utterly intent on our conversation. Why must I form any belief about the bird? Granted, if someone later asks if I saw a blue bird, I may assent, thereby indicating a belief that the bird *was* blue. But this belief is not perceptual: it is about a perceptible and indeed has visual content, but it is not grounded in seeing. Moreover, it may have been formed only when I recalled my visual experience of the bird. Recalling that experience in such a context may *produce* a belief *about* the thing I saw even if my original experience *of* the thing did not. For plainly a recollected sensory experience can produce beliefs about the object that caused it, especially when I have reason to gain information about that object. Perhaps one notices something in one’s recollected image of the bird, an image merely recorded in the original experience, but one formed no belief about the bird. Granted, perception must produce a sensory experience, such as an image, and granted such an image—and even a recollection of it—is raw material for beliefs; it does not follow that perception must produce beliefs.

It might be objected that genuinely seeing an object must produce beliefs, even if we are not *conscious* of its doing so. How else can perception guide our behavior, as it does when, on seeing a log in our path, we step over it?

One answer is that not everything we see, including the bird that flies by as I concentrate on something else, demands or even evokes a cognitive response, particularly one entailing belief-formation. If I am cataloguing local birds, the situation is different. But when an unobtrusive object we see—as opposed to one blocking our path—has no particular relation to what we are doing, perhaps our visual impressions of it are simply a *basis* for forming beliefs about it should the situation call for it, and it need not produce any belief if our concerns and the direction of our attention give the object no significance.

Despite the complexity I am pointing to in the relation between seeing and believing, clearly we may hold what is epistemologically most important here. Suppose I can see a bird without believing anything about (or of) it. Still, when I do see one, I *can* see it to be something or other, and my perceptual circumstances are such that I might readily both come to believe something about it *and* see that to be true of it. Imagine that someone suddenly interrupts a conversation to say, ‘Look at that bird!’ If I see it, I am in a position to form some belief about it, if only that it is swift, though I need not actually form any belief about it, at least not one I am conscious of.

To see these points more concretely, imagine I am alone and see the bird in the distance for just a second, mistakenly taking it to be a speck of ash. If there is not too much color distortion, I may still both know and justifiably believe it to be dark. Granted, I would misdescribe it, and I might falsely believe that it is a speck of ash. But I could still know something about it, and I might point the bird out under the misleading but true description, ‘that dark thing’. The bird *is* the thing I point at; and I can see, know, and justifiably believe that there is a dark thing there.

My perception of the bird, then, gives me a ready *basis* for some knowledge and justification, even if the perception occurs in a way that does not cause me to believe that there is, say, a *bird* before me and so does not give me *actual* knowledge of it. Seeing *is* virtual believing, or at least potential believing. A similar point holds for simple perception in the other senses, though some, such as smell, are in general less richly informative than sight.<sup>10</sup>

### *The perceptual hierarchy*

Our discussion seems to show that simple perceiving need not produce belief, and objectual perceiving need not always yield propositional perceiving. Still, this third kind of perception is clearly not possible without the first and, I think, the second as well. I certainly cannot see *that* the bird is anything if I do not see it at all; and I must also see it in order to see it *to be* something, say a speck of blue. Thus, simple perceiving is fundamental: it is required for objectual and propositional perceiving, yet does not clearly entail either. If, for instance, you do not perceive in the simple mode, say see a blue speck, you do not perceive in the other two modes either, say see a speck to be blue or see that it is blue. And as objectual perceiving seems possible without



propositional perceiving, but not conversely, the former seems basic relative to the latter.

### *Simple, objectual, and propositional perception*

We have, then, a perceptual hierarchy: propositional perceiving depends on objectual perceiving, which in turn depends on simple perceiving. Simple perceiving is basic, and it commonly yields, even if it need not always yield, objectual perceiving, which, in turn, commonly yields, even if it need not always yield, propositional perceiving. Simple perceiving, such as just seeing a green field, may apparently occur without either of the other two kinds, but seeing something *to be* anything at all, such as rectangular, requires seeing it; and seeing *that* it is something in particular, say green, requires both seeing it to be something and, of course, seeing it.

Thus, even if simple perception does not always produce at least one true belief, it characteristically does position us to form any number of true beliefs. It gives us cognitive *access* to perceptual information, perhaps even *records* that information in some sense, whether or not we register the information conceptually by forming perceptual beliefs of either kind.

### *The informational character of perception*

As this suggests, perception by its very nature is *informational*; it might even be understood as equivalent to a kind—a sensory kind—of receipt of information about the object perceived.<sup>11</sup> The point here is that not all perceptually given information is *propositional* or even conceptualized. This is why we do not receive or store all of it in the contents of our beliefs. Perceptual content—conceived as the content of a simple perception—is at least in part determined by the properties we are sensorily conscious of in having that experience; it is not equivalent to the content of the perceptual belief(s) that experience may produce.

Some of the information perception yields is imagistic. Indeed, we may think of all the senses as capable of yielding images or, for the non-visual senses, at least of yielding the non-visual counterparts of images—*percepts*, to use a technical term for such elements in perceptual experience occurring in any sensory mode, whether visual or auditory or of some other kind. It is in these sensory impressions that the bulk of perceptual information apparently resides. This point explains the plausibility of the idea that a picture is worth a thousand words—which is not to deny that, for some purposes, some words are worth a thousand pictures. A single report of smoke may avert a catastrophe; a single promise may alter a million lives.

It is in part because perception is so richly informative that it normally gives us not only imagistic information but also situational justification. Even if I could be so lost in conversation that I form no belief about the passing bird, I am, as I see it pass, normally justified in believing something about it,



concerning its perceptible properties, for instance that it glides.<sup>12</sup> There may perhaps be nothing highly specific that I am justified in believing about it, say that it is a cardinal or that its wingspan is ten inches, but if I really see it, as opposed to its merely causing in me a visual impression too indistinct to qualify me as seeing it, then there is something or other that I may justifiably believe about it.

When we have a clear perception of something, it is even easier to have perceptual justification for believing a proposition about it without actually believing it. Just by taking stock of the size of the field in clear view before me, I am justified in believing that it has more than 289 blades of grass; but I do not ordinarily believe—or disbelieve—any such thing about grassy fields I see. It was only when I sought a philosophical example about perception and belief, and then arbitrarily chose the proposition that the field has more than 289 blades of grass, that I came to believe this proposition. Again, I was justified in believing the proposition before I actually did believe it.

## Perceptual justification and perceptual knowledge

What is it that explains why seeing the bird or the field justifies us in believing something about what we see, that is, gives us situational justification for such a belief? And does the same thing explain why seeing something enables us to know various facts about it?

### *Seeing and seeing as*

One possible answer is that if we see something at all, say a bird, we see it *as* something, for instance black or large or swift, and we are justified in believing it to be what we see it as being. The idea is that all seeing and perhaps all perceiving is *aspectual perception* of a kind that confers justification. We see things by seeing their properties or aspects, for instance their colors or their front sides, and we are justified in taking them to have the properties or aspects we see them as having.

Let us not go too fast. Consider two points, one concerning the nature of seeing *as*, the other its relation to justification.

First, might not the sort of distinction we have observed between situational and belief justification apply to seeing itself? Specifically, might not my seeing the bird imply that I am only in a *position* to see it *as* something, and not that I *do* see it as something? It is true that when we see something, we see it *by* seeing some property or aspect of it; but it does not follow that we see it as *having* this property or aspect. I might see a van Gogh painting by its colors, shapes, and distinctive brush strokes, but not see it as having them because my visual experience is dominated by the painting as a whole. Someone might reply that if I see it by those properties, I am disposed to believe it has them and so must see it as having them; but this disposition

implies at most a readiness to see it as having them. There may, to be sure, be a sense in which if we see something aright, for example see a van Gogh with recognition of it as his, then we must see it as what we recognize it to be.

Seeing *as* can also be a matter of conceptualization—roughly, *conceiving as*. But this is different from perceptual seeing *as*. The distinction between perceptual seeing *as* and perceptual seeing *by* remains. Seeing *by* is causal and discriminative but not necessarily ascriptive or, especially, conceptual. Seeing *as*, though also causal, is often ascriptive and commonly conceptual. We see faces by seeing (for example) the distinctive shape of the eyes and mouth, but need not ascribe those to those we see or conceptualize these properties. But if we see a painting as blurry, we commonly ascribe that property to it and may conceptualize the painting as blurry.

Second, suppose that seeing the bird did imply (visually) seeing it *as* something. Clearly, this need not be something one is justified in believing it to be (and perhaps it need not be something one *does* believe it to be). Charles, our biased birdwatcher, might erroneously see a plainly black bird as blue, simply because he so loves birds of blue color and so dislikes black birds that (as he himself knows) his vision plays tricks on him when he is bird-watching. He might then not be justified in believing that the bird is blue.

Assume for the sake of argument that seeing implies seeing *as* and that typically, seeing *as* implies at least objectually believing something or other about the thing seen. Still, seeing an object as having a certain property—say, a stick in the water as bent—does not entail that it has the property. Nor does it always give one (overall) situational justification for believing it to have that property.

### *Perceptual content*

It is natural to think of perception as in some way *representational*. If we see things by seeing their properties, for instance, then our perceptual experience in some way represents the object as having them. If perceiving entailed believing, we could perhaps take it to have the same content of the entailed belief(s). But (simple) perception apparently does not entail believing, so this conception of its content is mistaken. For propositional and objectual perception, however, we might plausibly say something like this: the content of my perception that *p* includes both the proposition that *p* (hence also the content of that proposition) and also the content of my objectual perception of the thing in question; that content includes the properties I perceive the thing to have.

If we seek a broad notion of perceptual content for simple perception, we might say that all the properties represented in a perceptual experience constitute its content. Then, for greater specificity, we might call the totality of perceptually represented properties the *property content*. These include properties an object is seen *as* having.<sup>13</sup> They apparently also determine “what it

is like” to perceive the object, say a squirrel in a tree. In seeing it, one’s visual field is determined mainly by the grey, the distinctive furry shape, and the arboreal background.

For propositional and objectual perception, we might call the property-ascriptive propositions that the perceiver perceptually believes on the basis of the perceptual experience their *doxastic propositional content*. If we want to capture all the propositions that one might justifiedly believe (and know) on the basis of the perception, we might speak of its *total propositional content*. This would include such propositions as that the squirrel is crouching, has a nut in its mouth, is in sunlight, and many more that need not be believed as a result of simply seeing the animal.<sup>14</sup>

### *Seeing as and perceptual grounds of justification*

Whether or not seeing always implies seeing *as*, it does have property content and normally puts one in a position to form at least one justified belief about the object seen. Suppose I see the bird so briefly and distractedly that I do not see it as anything in particular; still, my visual impression of it has some feature or other by which I am justified in believing something of the bird, if only that it is a moving thing. Even Charles would be justified in believing something like this. His tendency to see black birds as blue is irrelevant to his perception of movement and does not affect his justification for believing such moving objects to be in motion.

Suppose, however, that for hours Charles had been hallucinating all manner of unreal things, and he knows this. Then he might not be justified in taking the bird he sees to be *anything* real, even though it is real. For as a rational person in this position he should see that if his belief is true, it may well be true only in the way a lucky guess is. Thus, the best conclusion here—and I suggest that this is an important justification principle concerning perception—is that normally, seeing an object gives one situational justification for believing something or other about it.

More broadly, it is very plausible to hold that *the evidence of the senses*—including above all the sensory experiences characteristic of perception—normally provides justification for beliefs with content appropriate to that evidence. If your experience is of a green expanse, you are justified in believing there is something green before you; if it is of something cool in your hand, you are justified in believing there is something cool in your hand; and so on.

One might also say something slightly different, in a terminology that is from some points of view preferable: seeing an object (always) gives one *prima facie justification* for believing something or other about it. *Prima facie* justification is roughly justification that prevails unless *defeated*. The two main kinds of defeater are such overriding factors as a strong justification for believing something to the contrary and such undermining (or undercutting) factors as my knowledge that I have been hallucinating and at present cannot

trust my senses. Overrides defeat *prima facie* justification by justifying an incompatible proposition instead; undermining defeaters simply prevent the would-be justification from succeeding. If I see a green field, I have a justification for believing it to be green; but I may not be justified, *overall*, in believing this if credible friends give me compelling reason to believe that despite appearances the field is entirely covered by blue grass, or that I am not seeing a field at all but hallucinating one.<sup>15</sup> In the former case, my justification is defeated by my acquiring better justification for a contrary proposition; in the latter, my visual justification is reduced below the threshold of success. If it is not eliminated, it is too weak to license saying I am justified in believing the proposition.

If seeing is typical of perception in (normally) putting us in a position to form at least one justified belief about the object seen, then perception in general normally gives us at least situational justification. This is roughly justification *for holding* a belief of the proposition for which we have the justification. As our examples show, however, it does not follow that every perceptual belief is justified. Far from it. Some perceptual beliefs, such as perceptual beliefs that are evidentially undermined by one's having formed similar beliefs based on hallucinations, are not. As with the biased bird-watcher, belief can be grounded in perception under conditions that prevent its being justified by that grounding.

Nevertheless, there is a simple principle of justification we can see to be plausible despite all these complexities: normally, a visual belief that is embedded in seeing that something is so or in seeing it *to be* so is justified (and it is always *prima facie* justified). If we see that an object has a property (say, that a field is rectangular) and, in virtue of seeing that it has that property (say, is rectangular), believe that it does, then (normally) we justifiedly believe that it does. Call this *the visual justification principle*, since it applies to cases of belief based on seeing that what is believed is true (or seeing it to be true).

I say *normally* (and that the justification is *prima facie*) because even here one's justification can be defeated. Thus, Charles might see that a bird is blue and believe on this basis that it is, yet realize that all morning he has been seeing black birds as dark blue and thus mistaking the black ones for the blue ones. Until he verifies his first impression, then, he does not justifiedly believe that the bird is blue, even though it in fact is. (We could say that he has *some* justification for believing this, yet better justification for not believing it; but to simplify matters I am ignoring degrees of justification.) He does indeed see a bird and may justifiedly believe that, but his belief that the bird is blue is not justified.

Suppose, on the other hand, that Charles has no idea that he has been hallucinating. Then, even when he does hallucinate a blue bird, he may be justified in believing that there is a blue bird before him. This suggests a related principle of justification, one that applies to visual experience whether it is a case of seeing or merely of visual hallucination: *When, on the basis of an apparently normal visual experience (such as the sort we have in seeing a bird*

*nearby*), one believes something of the kind the experience seems to show (for instance that the bird is blue), normally this belief is justified. Call this the *visual experience principle*, since it applies to cases in which one has a belief based on visual experience even if not an experience of actually seeing (the veridical kind). The visual principle takes us from seeing (vision) to justification; the visual experience principle takes us from visual experience—conceived as apparent seeing—to justification. The latter is wider: it indicates that visual experience can justify a huge range of beliefs, not just a belief to the effect that an object in fact has a property one sees it to have.

Similar principles can be formulated for all of the other senses, though the formulations will not be as natural. If, for example, you hear a note to be flat and on that basis believe that it is flat, normally your belief is justified. It is grounded in a veridical perception in which you have discriminated the flatness you believe the note has. And suppose, by contrast, that in what clearly seem to be everyday circumstances you have an utterly normal-seeming auditory hallucination of a flat note. If that experience makes it seem clear that you are hearing a flat note, then if you believe on the basis of the experience that this is a flat note, normally your belief would be justified. You have no reason to suspect hallucination, and the justification of your belief that the note is flat piggybacks, as it were, on the principle that normally applies to veridical beliefs.<sup>16</sup>

### *Seeing as a ground of perceptual knowledge*

Some of what holds for the justification of perceptual beliefs also applies to perceptual knowledge. Seeing the green field, for instance, normally yields knowledge about the field as well as justified belief about it. This suggests another visual principle, a visual knowledge principle. It might be called an *epistemic principle*, since it states a condition for the visual generation of knowledge: *At least normally, if we see that a thing (such as a field) has a property (say is rectangular), we (visually) know that it has it.* A parallel principle holds for objectual seeing: *At least normally, if I see something to have a property (say to be rectangular), I know it to have the property.*

There are, however, special circumstances that explain why these epistemic principles may have to be restricted to “normal” cases. It may be possible to see that something is so, believe on that basis that it is, and yet not know that it is. Charles’s case seems to show this. For if, in the kind of circumstances he is in, he often takes a black bird to be blue, then even if he sees that a certain blue bird is blue and, on that basis, believes it is blue, he apparently does not know that it is.<sup>17</sup> He might as well have been wrong, one wants to say; he is just lucky that this time his belief is true and he was not hallucinating. As he has no reason to think he has been hallucinating, and does not realize he has been, one cannot fault him for holding the belief that the bird is blue or regard the belief as inappropriate to his situation. Still, knowledge apparently needs better grounding than is provided by his blameless good fortune. This

kind of case has led some philosophers to maintain that when we know that something is so, our being right is not *accidental*.

There is an important difference here between knowledge and justification. Take knowledge first. If Charles is making errors like this, then even if he has no idea that he is and no reason to suspect he is, he does not know that the bird he believes to be blue is blue. But even if he has no idea that he is making errors, or any reason to suspect he is, he *may* still justifiedly believe that the bird is blue. The main difference between knowledge and justification here may be this: he can have a true belief that does not constitute knowledge because there is something wrong for which he is in no way criticizable (his errors might arise from a handicap which he has no reason to suspect, such as sudden color blindness); but he cannot have a true yet unjustified belief without being in some way criticizable. The standards for knowledge, one might say, permit fewer unsuspected weaknesses in discriminating the truth than those for justification, if the standards for knowledge permit any at all.

This difference between knowledge and justification must be reflected in the kinds of principles that indicate how justification, as opposed to knowledge, is generated. Justification principles need not imply that the relevant basis of a belief's justification assures its truth; but since a false belief cannot constitute knowledge, epistemic principles (knowledge principles) cannot capture elements that generate knowledge unless they rule out factors that might produce a false belief. A ground of knowledge must, in *some* way, suffice for the truth of the proposition known; a ground of justification must, in some way, *count toward* the truth of the proposition one is justified in believing, but need not rule out its falsehood.

On the basis of what we see, hear, feel, smell, and taste, we have a great many beliefs, propositional and objectual. There is apparently no good reason to doubt that these perceptual beliefs are commonly justified or that, quite often, they are true and constitute knowledge. But to see that perception is a basis of justification and knowledge is to go only part way toward understanding what perception, justification, and knowledge are. Here the main question is what constitutes perception, philosophically speaking. Until we have a good understanding of what it is, we cannot see in detail how perception grounds belief, justification, and knowledge. These problems cannot be fully resolved in this book, but we can achieve partial resolutions. I want to discuss (further) what perception is first and, later, to illustrate in new ways how it grounds what it does. The next chapter, then—also concentrating on vision—will start by considering some of the major theories of the nature of perception.

## Notes

- 1 Perceiving *of*, perceiving *to be*, and perceiving *that* may also be called perception of, perception to be, and perception that, respectively; but the

second expression is not common, and in that case at least, the *-ing* form usually better expresses what is intended.

- 2 A related way to see the difference between objectual and propositional beliefs is this. If I believe something to have a property, say a British Airways plane to be a Boeing 777, then this same belief can be ascribed to me using any correct description of that plane, say, as the most traveled plane in the British Airways fleet: to say I believe BA's most traveled plane to be a 777 is to ascribe the same belief to me. This holds even if I do not believe it meets that description—and it can hold even when I cannot understand the description, as a child who believes a tachistoscope to be making noise cannot understand 'tachistoscope'. By contrast, if I have a propositional belief, say that the United Airlines plane on the runway is the most traveled in its fleet, this ascription cannot be truly made using just any correct description of that plane, say the plane on which a baby was delivered on Christmas Day, 2001. I may have no inkling of that fact—or may mistakenly think it holds for a BA plane. A rough way to put part of the point here is to say that propositional beliefs about things are about them *under a description or name*, and objectual beliefs about things are not (even if the believer could describe them in terms of a property they are believed to have, such as being noisy). It is in part because we need not conceptualize things—as by thinking of them under a description—in order to have objectual beliefs about them that those beliefs are apparently more basic than propositional ones.
- 3 In terminology common in epistemology, objectual belief is *de re*—of the thing—whereas propositional belief is *de dicto*—of the proposition—and I am similarly distinguishing between objectual and propositional perception. The objectual cases, unlike the propositional ones, require no particular concept of the thing perceived. To be sure, those who do have the concept of a field and know that I believe it to be rectangular may say, 'He believes the field is rectangular', meaning that I believe it *to be* rectangular. English idiom is often permissive in this way, and in everyday life nothing need turn on the difference. Moreover, some philosophers have held that a thing, such as a field, can be a constituent in a proposition—in which case it might be considered a kind of content of a belief of that proposition—and this might provide a basis for saying that the two belief ascriptions may be properly interchangeable. I am ignoring that controversial and uncommon conception of a proposition. For detailed discussion of the extent to which perception is conceptual and of how it yields perceptual beliefs, see Michael Pendelbury, 'Sensibility and Understanding in Perceptual Judgments', *South African Journal of Philosophy* 18, 4 (1999), 356–69.
- 4 It may be best to leave open here that Susie could, at least for a moment, believe (in an admittedly weak sense of the term), of a tachistoscope, that it is making noise, yet not believe any proposition about it: she *attributes* noise-making to it, yet does not conceptualize it in the way required for



having a propositional belief about it, the kind of belief expressed in a complete declarative sentence such as 'The thing on the table is making noise'. She would then have no propositional belief about the instrument, the kind of belief that should unqualifiedly be called true (or false), such as that the tachistoscope is making noise. On this approach, what I am calling objectual belief is (or often is) better called *property attribution*. It is an attribution *to* the thing in question because of the kind of causal role that thing plays in grounding the attribution; and if it is not strictly speaking a belief, it does imply a disposition to form one, such as that the thing on the table is making noise.

- 5 Specifically, these are *doxastic* attitudes (from the Greek *doxa*, for 'belief'). A fear can be propositional and thereby cognitive, but it need not entail believing the proposition one fears is so, for example that the man approaching one will attack. Some might consider objectual awareness, say awareness of perfect symmetry, cognitive, at least when the person has the concept of relevant property. By contrast, desires, the paradigm *conative* attitudes, should not, I think, be taken to have propositional objects (e.g. 'to swim' in 'my desire to swim' does not express a truth or falsehood).
- 6 Perceptions that embody beliefs in the ways illustrated are also called *epistemic*, since the embedded belief is commonly considered to constitute knowledge. Their connection with knowledge is pursued in this chapter and others.
- 7 The distinction between simple and propositional perceiving and other distinctions drawn in this chapter are not always observed. At one point W.V. Quine says:

think of "x perceives y" rather in the image of "x perceives that p". We say "Tom perceives the bowl" because in emphasizing Tom's situation we fancy ourselves volunteering the observation sentence "Bowl" rather than "Surface of a bowl," "Front half of a bowl," "Bowl and background," and so on. When we ask "What did he perceive?" we are content with an answer of the form "He perceived that p".

(*Pursuit of Truth*, revised edn [Cambridge, MA: Harvard University Press, 1992], p. 65)

Notice that because seeing that (say) there is a bowl in front of one obviously entails seeing a bowl, it is no surprise that we are content with a report of the propositional perception even if we wanted to know only what object was seen: we get what we sought and more. It does not follow that simple seeing *is* or even entails propositional seeing. It is also worth noting that Quine is apparently thinking only of seeing here; for the other four senses, there is less plausibility in maintaining what he does.

- 8 The adage could not be taken to refer to simple seeing, for what we simply



see, say a glass or leaf or field, is not the sort of thing that can be believed (to be true or false). To be sure, seeing something, especially something as striking as golf ball-size hail, produces a *disposition to believe* certain propositions, say that this is a dangerous storm. But, by what seems an economy of nature, there are many things we are disposed to believe but do not. I have defended these points in detail in 'Dispositional Beliefs and Dispositions to Believe', *Noûs* 28 (1994), 419–34.

- 9 This applies even to full-scale philosophical dictionaries written by teams of experts, though such a work can provide concise statements of much valuable information. See, for example, the entries on blind sight and perception in Robert Audi (ed.), *The Cambridge Dictionary of Philosophy* (Cambridge: Cambridge University Press, 1995, 1999).
- 10 In the light of what has been said in this chapter so far we can accommodate much of what is plausible in the common view that, as D.M. Armstrong puts it:

[perception] is an acquiring of knowledge or belief about our physical environment (including our own body). It is a flow of information. In some cases it may be something less than the acquiring of knowledge or belief, as in the cases where perceptions are entirely discounted or where their content has been confidently anticipated.

(*Belief, Truth and Knowledge* [Cambridge: Cambridge University Press, 1973], p. 22)

First, I can agree that perception entails acquisition of *information*; the point is that *not all our information is possessed as the content of a belief*. Second, Armstrong himself notes an important way in which perception might fail to produce belief: it is “discounted,” as, for example, where one is sure one is hallucinating and so resolutely refuses to accept any of the relevant propositions.

- 11 This is the kind of view developed in detail by Fred Dretske. See esp. *Knowledge and the Flow of Information* (Cambridge, MA: MIT Press, 1981).
- 12 The notion of normality here is not statistical; it implies that what is not normal calls for explanation. In the world as we know it, exceptions to the normality generalizations I propose seem at least rare; but the point is not that statistical one, but to bring out that the very concepts in question, such as those of seeing and knowing, have a connection in virtue of which explanation is called for if what is normally the case does not occur.
- 13 A property that something is seen as having need not be a property it actually has; but here seeing *as* is phenomenal, not doxastic. Roughly, the perceptual content represents what the object is like if it in fact has the properties it is seen as having.

- 14 A detailed discussion of the representationality of perception and the kind of content it has is provided by Fred Dretske in *Naturalizing the Mind* (Cambridge, MA: MIT Press, 1997). He deals with the sense in which perceptual content is external. If, loosely speaking, we call the perceived object the *objectual content* then simple perception obviously has a kind of external content; but as the object is “in” the experience, it might be considered a kind of content, as indeed it may for propositional and objectual perceptions as well. With this idea in mind, it is clear how the perceptually believed propositions themselves may also be conceived as having external content. I have discussed internal and external content in relation to such examples in ‘Internalism and Externalism in Epistemology and Semantics’, in Mark Timmons, John Greco, and Alfred R. Mele (eds.), *Rationality and the Good: Critical Essays on the Ethics and Epistemology of Robert Audi* (Oxford: Oxford University Press, 2007). (This responds to a challenge from Timothy Williamson, ‘On Being Justified in One’s Head’, *ibid.*, 106–122)
- 15 In speaking of justification that prevails, and of overall justification, I have in mind the kind appropriate to a rational person’s believing the proposition in question, construed as roughly the kind such that when we believe a true proposition with that kind of justification then (apart from the kinds of case discussed in Chapter 10 that show how justified true beliefs *need not* constitute knowledge) we know it.
- 16 There are complexities I cannot go into, such as how one’s competence figures. I am imagining here someone competent to tell whether a note is flat (hence someone not tone deaf): in general, if we are not competent to tell whether a kind of thing has a property or not, an experience in which it seems to have it may not justify us in believing it does. There is also the question of *what* the belief is about when the “object” is hallucinatory, a problem discussed shortly. Still other problems raised by this justification principle are discussed in Chapter 11 in connection with the controversy between internalism and externalism.
- 17 If, as is arguable, seeing that it is blue entails knowing that it is, then he does *not* see that it is, though he sees its blue color. But this entailment claim is far from self-evident. Suppose he clearly sees a blue bird and believes it is blue, but does not know that it is because of his frequent hallucinations. A moment before, he hallucinated such a bird; a moment later, he will again; and he realizes his senses have been playing such tricks on him. Still, he cannot help believing this bird is blue and believes that *on the basis of* clearly seeing it and its color in normal light. Might we say that he sees that the bird is blue, but does not know this? We cannot say that he “can’t believe his own eyes,” because he does; but if, in the normal way, they show him the truth and he thereby believes it, might he not see it through them?



## 2 Theories of perception

Sense experience, appearances, and reality

- **Some commonsense views of perception**

*Perception as a causal relation and its four main elements*

*Illusion and hallucination*

- **The theory of appearing**

- **Sense-datum theories of perception**

*The argument from hallucination*

*Sense-datum theory as an indirect, representative realism*

*Appraisal of the sense-datum approach*

- **Adverbial theories of perception**

- **Adverbial and sense-datum theories of sensory experience**

- **Phenomenalism**

*A sense-datum version of phenomenalism*

*Adverbial phenomenalism*

*Appraisal of phenomenalism*

- **Perception and the senses**

*Indirect seeing and delayed perception*

*Sight and light*

*Vision and the eyes*

## 2 Theories of perception

### Sense experience, appearances, and reality

Much has now been said about perception and its relation to belief, justification, and knowledge. What we have seen puts us in a position to frame a theory of what perception is, but we have not yet stated any such theory. Stating and assessing theories of perception is the task of this chapter.

#### Some commonsense views of perception

One natural thing to say about what it is for us to see the green field is appealingly brief. We simply see it, in an ordinary way: it is near and squarely before us; we need no light to penetrate a haze or a telescope to magnify our view. We simply see the field, and it may normally be taken to be pretty much as it appears. This sort of view, called *naïve realism*, has been thought to represent common sense: it says roughly that perception is simply a matter of the senses telling us about real things.

The view is naïve because it ignores problems of a kind to be described in a moment; it is a form of realism because it takes the objects of perception to be real things external to the perceiver, the sorts of things that are “out there” to be seen whether anyone sees them or not.

A more thoughtful commonsense view retains the realism without the naivety. It is quite commonsensical, for instance, to say that I see the field *because* it is before my open eyes and stimulates my vision, thereby *appearing* to me as a green, rectangular shape. Stimulating my vision is a causal relation. For instance, the field, by reflecting light, causes me to have the visual experience that is part of my seeing that very field. Moreover, the field apparently must cause my visual experience if I am to *see* it. As the more thoughtful commonsense view specifies that the object of perception must be a real external thing, we might call it a *perceptual realism*. Most—but not all—theories of perception incorporate perceptual realism.

To understand why perception must have a causal element, suppose I am looking at the field and, without my noticing, someone instantaneously drops a perfect picture of the field before my eyes. My visual experience might not change. What appears to me might look just as the field did. Yet I no longer *see* the field. Instead, I see a picture of it. (I do see the field *in* the picture, but

that is secondary seeing and not the kind we are talking about.) The reason I do not now see the field is roughly that it has no (causal) effect on my visual experience.

### *Perception as a causal relation and its four main elements*

Examples like this suggest that *perception is a kind of causal relation* between whatever is perceived and its perceiver, wherein the object perceived produces a sensory experience in the perceiver. This is a plausible, commonsensical, and important point, though it does not tell us precisely what perception is. I call any theory of perception which incorporates the point *a causal theory of perception*. Most theories of perception are causal.

We can now better understand the four elements I have described as among those crucial in perception: the perceiver, the object perceived, the sensory experience in which the object appears to the perceiver, and the causal relation between the object and the perceiver, the relation wherein the object produces that experience. If you see the field, there is a distinctive way, presumably involving light transmission to your eyes, in which the field produces the visual sensory experience of a green, rectangular shape characteristic of your seeing it. If a picture of the field produces an exactly similar visual experience in the same way, it is the picture you see, not the field. Similarly, if you hear a piano piece, there is a special way in which it causes you to have the auditory sensations of chords and melody and harmony that go with it.

It is difficult, though fortunately not necessary for a general understanding of perception, to specify precisely what these causal paths from the object to the perceiver are. Some of the details are the business of the psychology and neurophysiology of perception. Others are determinable by philosophical inquiry. Philosophical reflection shows us, for instance, that not just any causal chain is the right sort for perception. Consider what is sometime called a *wayward (or deviant) causal chain*. Suppose the piano sounds cause a special machine, created by a prankster, to produce in me both temporary deafness and a faithful auditory hallucination of the piece. Then I do not *hear* it, though my sensory experience, the auditory experience I enjoy in my own consciousness, is just what it would be if I did hear it. Nor do I hear it if, though the sound waves reach my brain and cause me to believe a piano is playing just the piece in question, I have no auditory experience. Even such a highly informed inner silence is not musical.

### *Illusion and hallucination*

We can make progress by pursuing the question of why naive realism is naive. Suppose there is a gray haze that makes the green field look gray. Or suppose the mouth of the cup I am holding appears, from a certain angle, as if it were an ellipse rather than a circle, or feels warm only because my hand is cold.

These are *perceptual illusions*, roughly perceptual experiences that (in the way illustrated) misrepresent the object of perception. They illustrate that things are not always as they seem. The cup is round and at room temperature.

Now imagine that the field burns up. I sorely miss its rich green and the spruce and maple, and on waking from a slumber in my chair I have a *hallucination*, roughly a sense experience qualitatively like a perceptual one but not of an external object. In this case, my (hallucinatory) visual experience is just as it would be if I were seeing the field. Here the grass I seem to see is not there at all. The point is not that something I see is not as it seems (as in the case of illusion) but that there seems to be something where there is nothing. With illusion, as illustrated by a partly submerged stick's looking bent, what is there is perceived distortedly; with hallucination, it appears that something is perceived when nothing is. Illusions and hallucinations are possible for the other senses too. When they occur, we do not just see (or hear, taste, smell, or touch) the object. Either we do not see it as it is or (perhaps) we do not see anything at all. Not everything we perceive is as it appears to be, and naive realism does not explain why.

One way to deal with illusion and hallucination is to stress how they show the need to distinguish appearance from reality. In a visual illusion, one sees something, but it does not appear as it really is, say circular. In a hallucination, nothing need appear to me at all, and if anything does, it is in reality even less what it appears to be than is the object of an illusion, or is not what it appears to be at all: instead of a blue spruce tree's appearing blue to me, for instance, it is as if the conical section of space where it stood appears "bespruced."

## The theory of appearing

The sort of account of perception just sketched as an improvement over naive realism has been called *the theory of appearing*: it says roughly that perceiving an object, such as a book, is simply its appearing to one in a certain (sensory) way. It may, for instance, appear to be rectangular. Thus, one perceives it—in this case, sees it—as rectangular. The theory can also provide the basis of an account of the sort of experience we have in hallucination as opposed to normal perception: that experience, too, the theory takes to be a case of something's appearing to one to have a set of properties; the object that appears is simply a different kind.<sup>1</sup>

The theory of appearing is initially plausible. It incorporates much reflective common sense, for instance the view that if one sees something, then it appears to one in some way, say as a red barn or at least as a visually experienced rectangular patch. The theory says nothing, however, about the need for a causal relation between the object and its perceiver (though it allows that there be one). If, consistently with its commonsense motivation, one stipulated that the crucial relation of appearing to the perceiver to have a

property—say to be rectangular—is a causal relation, one would then have a different theory (of a kind to be discussed shortly).

In addition to the question of how the theory can do justice to the causal element in perception, it faces a problem in accounting for hallucinations in which there is no object to appear to the person at all. I could, after all, hallucinate a green field when I see nothing, say because it is pitch dark or my eyes are closed. In such an *empty hallucination*—one that occurs despite my perceiving nothing—what is it that appears green to me? One plausible answer is given by a quite different theory of perception. Let us explore that view.

### **Sense-datum theories of perception**

Once we think seriously about illusion and hallucination, we may begin to question not only naive realism but also any kind of *direct realism* in the theory of perception, any perceptual realism which, like the theory of appearing, says that we see (or otherwise perceive) external objects directly, rather than *through* seeing (or at least visually experiencing) something else. After all, not only do light rays come between us and what we see, there are also brain events crucial for seeing. Perhaps these events or other intermediaries in perception produce or indicate an interior object, presumably a mental object that plays an intermediary role in perception.

Hallucination illustrates most readily how such an intermediary might figure in understanding perception. Imagine that you vividly hallucinate the field just as it would be if it were before you. This seems quite possible. If such a “faithful” hallucination occurs, your visual experience—roughly, what you are aware of in your visual consciousness—is exactly like the experience you have when you see the field. Does it not then seem that the difference between ordinary seeing and visual hallucination is simply in what *causes* the visual experience, rather than in what you directly see? When I see the field, *it* causes my visual experience. When I hallucinate it, something else (such as my deep desire to have it back) causes my visual experience. But apparently what I directly see—the immediate object of my visual experience—is the same object in both cases. This point presumably explains why my visual experience is qualitatively the same whether I am hallucinating the field or really seeing it. If it were not the same, we could not say things like ‘It was exactly as if I were seeing the tree in normal light’.

### ***The argument from hallucination***

We might develop these ideas by considering an argument from hallucination. It consists of two connected arguments. The first constituent argument attempts to show a parallel between hallucination and ordinary perception:



- 1 A perfectly faithful (visual) hallucination of a field is qualitatively just like, and so intrinsically indistinguishable from, an ordinary experience of seeing that field, that is, not distinguishable from it just in itself *as* a visual experience, as opposed to being distinguishable through verifying one's visual impression by touching the things around one.<sup>2</sup>

Hence:

- 2 What is *directly* "seen," the immediate object of one's visual experience, is the same sort of (non-physical) thing in a perfect hallucination of a field as in an ordinary experience of seeing a field.

But—and we now come to the second constituent argument, which builds on (2) as its first premise—clearly:

- 3 What is directly seen in a hallucination of a field is not a field (or any other physical thing).

Indeed, no field is seen at all in a hallucinatory visual experience, so (3) seems plainly true. Hence, putting (1)–(3) together, we may infer that:

- 4 What is directly seen in an ordinary experience of seeing a field is not a field.

The overall idea is that when we ordinarily see an everyday perceptible object such as a field, we see it through seeing—in the sense of visually experiencing—something else *directly*: something not seen *by* seeing anything else. What we see directly—call it a *sense-datum*—might be an image. One may prefer (as some philosophers do) to say that we do not in any sense *see* such things but are only visually acquainted with them. To simplify, let us bear this alternative in mind but use the more natural term 'see'.

Just what is directly seen when one sees a field, then, and how is the field *indirectly* seen? Why not say that what is *directly* seen is a two-dimensional object consisting of the same sorts of colors and shapes one sees in the hallucinatory experience? After all, nothing, not even (physical) light, intervenes between us and them. There is no "space" for intermediaries. Hence, no intermediaries can misrepresent these special objects. These objects are apparently internal to us: as traditionally conceived, they could exist even if—as Descartes, in his *Meditation I*, supposed to be possible—we were disembodied minds in an otherwise empty world. The only space they need is in the mind. Yet we do see the field *by* seeing them; hence, we see it indirectly, at least in the sense that we see it by having a visual experience that is directly of something else.

The idea that experiencing sense-data is required for perception is nicely expressed in Emily Dickinson's poem 'I Heard a Fly Buzz When I Died'. In the final moment of her terminal experience:

There interposed a fly,  
 With blue, uncertain stumbling buzz,  
 Between the light and me;  
 And then the windows failed, and then  
 I could not see to see.

The external light from the window blocks her eyesight, but this leaves inner seeing—portrayed here as a necessary condition for ordinary seeing—still possible. Until the end, she can see *to* see. It is sense-data that are conceived as the direct objects of such inner sight.

A sense-datum theory is perfectly consistent with a causal theory of perception: the field produces the colors and shapes in my visual consciousness in a way that fully accords with the view that perception is a causal relation between something external and the perceiver. Perception is simply a *mediated*, hence *indirect*, causal relation between external objects we perceive and us: the object produces the mediating colors and shapes that appear in our visual fields, and, through seeing them, we see it.

The theory I am describing is a version of a *sense-datum theory of perception*. Such theories are so called because they account for perception by appeal to a view of what is directly given *in* sense experience, hence is a *datum*, a given, for such experience—the sort of thing one is visually aware of in hallucinating a field. This sense-datum thesis (unlike the phenomenalist sense-datum view to be discussed shortly) is a realist view; but its realism, by contrast with that of naive realism and the theory of appearing, is indirect.<sup>3</sup>

### *Sense-datum theory as an indirect, representative realism*

A sense-datum theory is a kind of *representative realism* because it conceives perception as a relation in which sense-data represent perceived external (hence real) objects to us. On some conceptions of sense-data, they are copies of those objects: shape for shape, color for color, sound for sound. John Locke held a view of this kind (and in 1689 published it in *An Essay Concerning Human Understanding*, especially Books II and IV), though for him sense-data are copies (“resemblances”) only of the *primary qualities* of physical things—solidity, extension (in space), shape, and mobility—not of their *secondary qualities*, above all colors, sounds, smells, and tastes. (He took the primary qualities to be objective and of the kind that concern physical science; and he considered the secondary ones to be in a sense subjective, not belonging to physical things but something like representational mental elements that they cause in us. Color, then, disappears in the dark, though the physical object causing us to see it is not changed by the absence of light.) Our question is whether any sense-datum version of representationalism is sound, and we need not pursue the interesting question of how these two kinds of qualities differ.

Sense-datum theories have had illustrious defenders down to the present

age. The theory has also had powerful opponents. To appreciate it better, let us first consider how it takes perception to be indirect. Sense-datum theorists might offer several reasons to explain why we do not ordinarily notice the indirectness of perception (I speak generally, not solely of Locke's version of the theory). Here are two important reasons.

First, normally what we directly see, say colors and shapes, roughly corresponds to the physical objects we indirectly see by means of what we see directly. It is only when there is an illusion or hallucination that we are forced to notice a discrepancy between what we directly see and the object commonly said to be seen, such as a book.

Second, the beliefs we form on the basis of perception are formed spontaneously, not through any process requiring us to consider sense-data. Above all, we do not normally *infer* what we believe about external objects we see from what we believe about the colors and shapes we directly see. This is why it is easy to think we “just see” things, directly. Perceiving is not inferential, and for that reason (perhaps among others) it is not *epistemically indirect*, in the sense that *knowledge* of external objects or beliefs about them are indirect, in the sense that they are based on knowledge of sense-data, or beliefs about them. On a plausible sense-datum view, I know that the field is green through *having* rectangular green sense-data, not through *inference from* propositions about them.<sup>4</sup>

It is apparently true that, as a sense-datum view may allow, perception is not inferential or epistemically indirect in the way inferentiality would imply. But, for sense-datum theorists, perception is nonetheless causally and objectually indirect. The perceived object is presented to us via another object, though not by way of a *premise*. These theories are *causally indirect*, then, because they take perceived physical objects to cause sensory experience, say of colors and shapes, *by* causing the occurrence of sense-data, with which we are directly (and presumably non-causally) acquainted in perceptual experience. Perception is also *objectually indirect* because we perceive external things, such as fields, *through* our acquaintance with other objects, namely sense-data. Roughly, we perceive external things through perceptual acquaintance with internal things.

Despite the indirectness of perception in these two respects, a sense-datum theorist need not deny that we normally do not use information about sense-data to arrive at perceptual beliefs inferentially, say by an inference from my directly seeing a grassy, green rectangular expanse to the conclusion that a green field is before me. Ordinarily, when I look around, I form beliefs about the external environment and none at all about my sensory experience. That experience causes my perceptual beliefs, but what they are *about* is the external things I perceive. It is when the colors and shapes do not correspond to the external object, as when a circle appears elliptical, that it seems we can understand our experience only if we suppose that the direct objects of sensory experience are internal and need not match their external, indirect objects. Is the sense-datum view, however, correct on this point?

### *Appraisal of the sense-datum approach*

Let us focus first of all on the argument from hallucination, whose conclusion suggests that what is directly seen in visual perception of external objects are sense-data. Suppose I do have a hallucination that is qualitatively just like, and intrinsically indistinguishable from, the normal experience of seeing a field. Does it follow that what is directly seen in the hallucination is the same sort of thing as what is directly seen in the normal experience? At least two problems confront the sense-datum theory here.

First, why must *anything* be seen in a hallucination? Imagine that I hallucinate the burned-up field. I might get up, still half asleep, and, pointing to the area, cry out, "It's regrown!" You might conclude that I *think* I see the field again. My *initial* reaction to realizing I had hallucinated the field might be that, hallucination or no, I *saw* it. But I might just as easily slump back in my chair and mumble that I wish I had seen it.

A compromise view would be that I saw the hallucinated grass *in my mind's eye*. But suppose I did see it in my mind's eye, and again suppose that the hallucination is intrinsically just like the ordinary seeing. Does it follow that what I directly see in the ordinary experience is the same as what I see in the hallucination, namely, something in my mind's eye? It does not. The notion of seeing in one's mind's eye is metaphorical, and such seeing need not imply that there is any real object seen, in or outside the mind. However vividly I may, in my mind's eye, see myself standing atop a giant pyramid in Toronto, that city has no pyramid, nor need there be any pyramidal object in my mind.

There is a second reason to resist concluding that something must be directly seen in hallucinations. Recall that my seeing a green field is apparently a causal relation between a sensory experience in me and the field that produces the experience. If so, why should the possibility that a hallucination can mimic my seeing the field tell us anything about what is directly seen (or is an object of visual acquaintance) when one sees that field? It is not as if we had to assume that only an experienced *object* can produce, or causally sustain, the relevant sensory experience, and must then conclude that it is an internal perceptual object, since there is no other candidate. Many things can have more than one causal basis, and the sense-datum theorist has no argument to show that only an internal perceptual *object*, or an acquaintance with it, as opposed, say, to an abnormality in the visual cortex (which need not be an object at all), can cause or sustain the hallucinatory experience.

Moreover, from the similarity of the internal, experiential elements in the hallucination and the internal ones occurring in genuine perception, one might as well conclude that since the ordinary experience is one of seeing only an external rather than an internal object, the hallucinatory experience is different only in the absence of the external object. Rather than add to the components that seem needed to account for the ordinary experience, we

subtract one that seems needed to account for the hallucination. This yields a more economical theory of perception.

An analogy may help. Compare trying to infer facts about how we see an original painting from facts about how we see it in a photo of it. From the indirectness of the latter seeing, it certainly does not follow that ordinary seeing of the painting is indirect. And even if a photographic viewing can be so realistic that it perfectly mimics an ordinary viewing, it does not follow that photographic, two-dimensional objects are components in ordinary seeing. Similarly, no matter how much like ordinary experiences hallucinations can seem, it does not follow that the former have all the internal elements (roughly, mental or mind-dependent elements) of the latter.

It will help to consider a different analogy. Two perfect ball-bearings, by virtue of identical diameter and constitution, can be qualitatively alike and so intrinsically indistinguishable, yet they can still differ, one being on your left and one on your right. Their intrinsic (i.e., non-relational) properties can thus be identical, while their *relations* (to you) differ. Hence they *do* differ in their relational properties. Similarly, the hallucination of a field and the ordinary visual experience of a field can be qualitatively alike, and so intrinsically indistinguishable, yet differ in their relations to me or to other things. One of them, the visual experience of a field, may be an element in a perceptual relation to the field; and the experience we call hallucination, which is not based on perceiving the external object hallucinated, may *not* be an element in any perceptual relation to the field, but only something I undergo (an element simply “in” me, on the plausible assumption that it is mental).

To account for the difference between the two kinds of experience, we might say this: the visual experience indicates the presence of an external thing; the hallucinatory experience, though intrinsically just like the visual one, only pretends to indicate an external thing to me. Thus, for all the argument from hallucination shows, the ordinary experience of seeing might be a relation to an object such as a green field, namely the relation of directly seeing, while the hallucinatory experience of a green field is not a relation to that field, such as being an internal copy of it, or even a relation to any other object, such as a perceiver.

The points just made about the argument from hallucination indicate that it is not sound. Its first premise, (1), does not entail (2) the conclusion drawn from it. Nonetheless, the argument poses serious problems for alternative theories. What explanatory account of hallucinations and illusions besides the sense-datum account might we adopt? To see some of our alternatives, it is best to begin with illusion rather than hallucination.

Recall the mouth of the cup viewed from an angle. A sense-datum theory will say we directly see (or anyway experience) an elliptical shape and indirectly see the cup. The theory of appearing, however, can also explain this: it reminds us that things need not be what they appear to be and says simply that they can appear elliptical even if they are round.

## Adverbial theories of perception

One could also combine the causal element in the sense-datum approach with the direct realism of the theory of appearing and move to a third theory, one that says the cup causes us to see it directly, rather than through producing sense-data in us, yet (because of our angle of vision) we see it as if it were elliptical. To avoid suggesting that anything in one's experience need *be* elliptical, one could take this to mean that the cup visually "appears elliptically" to us. Here the adverb 'elliptically' modifies the verb 'appears' and describes a *way* in which we visually experience the cup. It does not imply that there is an object that appears to us and *is* elliptical.<sup>5</sup> Let us explore this idea in relation to the theory associated with it.

It should now be clear why we need not grant (what sense-datum theorists sometimes seem to assume about perception) that in order for an object to appear a given way to us there must *be* something we see that *is* that way, for instance an elliptical sense-datum. Suppose that one says simply that the cup appears elliptically, using this adverb to designate, from the perceiver's point of view, *how* one visually experiences it: elliptically. To say it appears elliptically is roughly to say it appears in the way an ellipse does, viewed from directly above its center, as opposed to the way a circle does when so viewed.

If this adverbial interpretation of such statements as 'I see an ellipse' seems artificial, consider an ordinary analogy. If I say I have a fever, no one could plausibly insist that there is an object, a fever, which I have. 'I have a fever' is a way of saying I am feverish, that is, my body is above a certain temperature. What our language seems to treat as a statement of a relation to an object, a fever, is really an ascription of a property: the property of having a temperature above a certain level. Just as 'having a fever' can ascribe a certain temperature, 'seeing elliptically' (in illusional and hallucinatory cases) can ascribe a certain visual experience.

On the basis of this move, one can construct what is called *the adverbial theory of perception*. Unlike the theory of appearing, which takes perception to be an unanalyzable relation in which things appear to us as having one or more properties, an adverbial theory conceives perception as an analyzable way of experiencing things. In what may be its most plausible form, it says roughly that *to perceive an object is for that object (in a certain way) to produce in one a sensory experience of it*: to cause one's experiencing it in a certain qualitative way, say to see a stick as straight (or, given the illusion induced by partial submersion, as bent). Both theories are, however, direct realist views. In both, the perceptual object appears to the perceiver without an intermediary. Other similarities (and some differences) between the two theories will soon be apparent.<sup>6</sup>

The adverbial theorist stresses that we see (or otherwise perceive) things in a particular qualitative way and that they thus appear to us in that way. Often they appear as they are; sometimes they do not. In each case they

are seen directly, not through intermediaries. Even if I do not see the cup as circular, I do see *it*: it is seen directly, yet appears elliptically.

So far, so good, perhaps. But what about hallucinations? Here the adverbial theory again differs from the theory of appearing. Unlike the latter, it denies that all sensory experience is *of* some object. The importance of this denial is not immediately apparent, perhaps because we suppose that usually a person visually hallucinating does see *something*. Consider Shakespeare's Macbeth, distraught by his murder of Duncan, hallucinating a dagger that seems to him to hover in mid-air:

Is this a dagger which I see before me,  
The handle toward my hand? Come, let me clutch thee.  
I have thee not, and yet I see thee still.  
Art thou not, fatal vision, sensible  
To feeling as to sight? or art thou but  
A dagger of the mind, a false creation,  
Proceeding from the heat-oppressed brain?  
I see thee yet, in a form as palpable  
As this [sword] which now I draw.

(Act II, scene i)

Perhaps Macbeth does see something in the place in question, say the wall behind "the dagger" or at least a chunk of space where it hovers. Thus, to explain what it is for Macbeth to experience "daggerly," an adverbial theorist might posit an "object" where the "dagger" seems located, if only the section of space where it seems to be. On one view, this object might be thought to play a role in causing him to have daggerish visual sensations. But there need not be any such object in a hallucination. For the theory of appearing, too, the space before him, despite being transparent, might somehow appear to him to be a dagger.

Supposing we accept this adverbialist account, what happens if it is pitch dark and Macbeth's hallucination is therefore *empty*, in the sense that there is nothing he sees, and hence nothing to serve as an object visually distorted into an apparent dagger? Then, whereas the theory of appearing may have to posit a sense-datum (or other special kind of object) to serve as what appears to be a dagger, the adverbial theory can deny that there is *any* kind of object appearing to Macbeth. It may posit some quite different account of his "be-daggered" visual experience, such as a psychological account appealing to the influence of drugs or of his "heat-oppressed brain."

Is it really plausible to hold, with the adverbial theory, that Macbeth saw nothing at all? Can we really explain how the normal and hallucinatory experiences are intrinsically alike without assuming they have the same direct objects? In the light of the special case of empty hallucination, the sense-datum theory may seem the most plausible of the three. It provides an object of Macbeth's visual experience even if that occurs in utter darkness,

whereas the adverbial theory posits no perceptual objects at all in empty hallucinations. Moreover, the sense-datum view postulates the same sort of direct object for ordinary perception, illusion, and hallucination, whereas the theory of appearing does not offer a uniform account of their direct objects and must explain why entities such as sense-data do not occur in normal perception as well as in empty hallucination.

Perhaps, however, the hallucination problem seems more threatening than it should to the adverbial theory because hallucinations are felt to be *perceptual* experiences and hence expected to be *of* some object. But as we have seen, although hallucinatory experiences can be intrinsically indistinguishable from perceptual ones, all that can be assumed is that the former are *sensory experiences*. Hallucinatory experiences, on the adverbial view, are simply not cases of perceiving, at least not in a sense requiring that any object appear to one.

Thus, nothing at all need appear to one in hallucinations, though it may *appear to the subject* that something is there. The hallucinator may then be described as having a visual sensory experience, but—as nothing is perceived—not a genuine perceptual experience.

### **Adverbial and sense-datum theories of sensory experience**

A perceptual experience is always sensory, and normally a sensory experience of the sort we have in perceiving is genuinely perceptual. But, as hallucination shows, a kind of short-circuit can cause the sense-receptors to produce sensory experience that is not a normal perceptual experience (or even part of one). It is important to consider the debate between adverbial and sense-datum theories in relation to sensory experience. Both theories take such experience to be essential to perception, and both offer accounts of sensory experience as well as of perception.

The most natural thing for adverbial theorists to say about hallucinatory experience is that it is not genuinely perceptual, but only sensory. They might, however, say instead that when a perceptual experience is hallucinatory, it is not a case of *seeing* (except perhaps in the mind's eye, or perhaps in the sense that it is seeing colors and shapes conceived abstractly as properties and not as belonging to sense-datum objects). The former description accords better with how seeing is normally understood: normally, we cannot be said to see what is not there.

The theory suggested by these responses to the hallucination problem might be called *the adverbial theory of sensory experience*. It says that having a sensory experience, such as a hallucination of a green field, is experiencing in a certain *way*, for example visually experiencing “green-fieldly.” Our commonsense assumption is that hallucination is not usual (for normal people) and that most of our vivid sensory experiences are genuinely perceptual. They are of, and thus caused by, the external object(s) apparently perceived.



But some sensory experiences are neither genuinely perceptual nor externally caused. People having them are in, say, a vision-like state, and what occurs in their visual cortex may be the same sort of process that occurs when they see things. Yet they are not seeing, and their visual experience typically has an internal cause, such as an abnormal emotion.

May we, then, regard sense-datum theories of perception as refuted by the points just made in criticism of the argument from hallucination and on behalf of the suggested adverbial theory and the theory of appearing? Certainly not. We have at most seen how one major argument for a sense-datum theory of perception fails and how alternative theories of perception can account for the apparently central elements in perception: the perceiver, the (ordinary) object perceived, the sensory experience, and the causal relation between the second and third.

Indeed, supposing that the argument from hallucination fails to show that sense-data are elements in normal everyday perception, sense-data might still be needed to account for non-perceptual sensory experience (sometimes loosely called perceptual experience because it is characteristic of perception). In this limited role, one might posit a *sense-datum theory of non-perceptual sensory experience*: such experience is considered direct acquaintance with sense-data.

A sense-datum view may seem preferable to an adverbial theory of sensory experience. For one thing, there is something unsatisfying about the idea that even in a visual hallucination so vivid that, if one did not suspect error, one would stake one's life on the presence of the hallucinated object, one sees nothing, except either in one's mind's eye, or in a sense of 'see' which does not require that any object be seen. Still, perhaps there is such a sense of 'see'.

There is another aspect of the controversy. It concerns the *metaphysics* associated with adverbial and sense-datum theories of any kind, specifically the sorts of things they require us to take as fundamental realities. In this respect, the adverbial theories of perception and sensory experience have an advantage over the counterpart sense-datum theories: the former do not posit a *kind* of object we would not otherwise have to regard as real. From the adverbial perspective, the objects that perception and sensory experience involve are simply perceivers and what they perceive. These are quite familiar entities which we must recognize and deal with anyway.

Sense-data are quite different from ordinary (presumably physical) objects of perception. Sense-data are either mental or at least depend for their existence on the mind of the subject. Yet they are unlike some mental phenomena in that no plausible case can be made for their being really brain phenomena, since they have properties, for instance being green and perfectly rectangular, not normally found in the brain.<sup>7</sup>

Moreover, there are difficulties in fully understanding sense-data in any terms. Is there, for instance, even a reasonable way of counting them? Suppose my image of the green field gradually gets greener. Is this a sense-datum changing or a new one replacing an old one? There seems to be no

way to tell. If there is no way to tell, how can we ever be sure we learn more about a sense-datum than what initially appears to us in experiencing it: how can one distinguish learning something more about *it* from learning about something new?<sup>8</sup>

Problems like these also affect the theory of appearing insofar as it must posit sense-data or similar entities to account for hallucinations. To be sure, such problems can also beset our understanding of ordinary objects. Can we always distinguish a mountain with two peaks from two mountains, or one snarled barberry bush from two with intertwined roots? But these problems seem less serious, if only because there is no question that there are *some* things of the physical kind in question. The corresponding problems may in the end be soluble for sense-data, but they at least give us some reason to prefer a theory that does not force us to regard sense-data as the only objects, or as even among the objects, we are directly aware of when we see, hear, touch, taste, and smell.

## Phenomenalism

If some philosophers have thought that perception can be understood without appeal to sense-data, others have conceived it as understandable in terms of sense-data alone as its objects. This view has the advantage of being, in at least one way, simpler than the adverbial and sense-datum theories. But the view is motivated by other considerations as well.

### *A sense-datum version of phenomenalism*

Think about the book you see. It is a perceptible object. Suppose we may conceive a real perceptible object as a perceptible object that is as it is, independently of what we think it to be. Still, real perceptibles, such as tables and chairs and books, are also plausibly conceived to be, by their very nature, *knowable*, at least in being experienceable. Indeed, it is doubtful that real objects of this sort could be unknowable, or even unknowable through the senses if lighting and other perceptual conditions are good. Now suppose we add to these ideas the assumption that our only genuine, certain knowledge of perceptibles is restricted to what directly appears to us and would be as it is even if we should be hallucinating. And what more does appear to us besides colors, shapes and other sensory (sensible) properties? Further, how do we know that this book, for example, could even exist without someone's perceiving its sensory properties? Certainly we cannot *observe* the book existing unperceived. If you observe it, you perceive it.

Moreover, if you imagine subtracting the book's sensory properties one by one—its color, shape, weight, and so on—what is left of it? This is not like peeling an apple, leaving its substance. It is like stripping layer after layer from an onion until nothing remains. Might we not conclude, then, that the book is not only *known by* its sensory properties, as the other theories of

perception also hold, but also *constituted by* a stable collection of them—by visual, tactual, and other sense-data that recur in our experience, confronting us each time we have the sense-data corresponding to, say, a certain bookcase in our home? Similarly, might it not be that to see the book is simply to be visually acquainted with such a stable collection of sense-data?

George Berkeley argued from a variety of angles that this is indeed what a perceptible object is. This view (which Berkeley developed in detail in his *Treatise Concerning the Principles of Human Knowledge*, published in 1710) is a version of what is often called *phenomenalism*, as it constructs external objects out of phenomena, which, in this use of the term, are equivalent to sense-data. The view is also considered a kind of *idealism*, since it construes physical objects as ideal, in the sense of being composed of “ideas” (an old term for sense-data) rather than material stuff that would exist even if there were no minds and no ideas.<sup>9</sup>

### *Adverbial phenomenalism*

Phenomenalism as just described is focused on the nature of perceptible objects but implies a related view of perception. In the sense-datum version of phenomenalism we have been examining, the associated account of perception retains a sense-datum theory of sensory experience, but not a sense-datum theory of perception. The latter view posits external objects as causes of the sense-data experienced in ordinary perception, whereas sense-datum phenomenalism says that physical objects *are* collections of sense-data.

Using the adverbial theory of sensory experience, one might also formulate an *adverbial phenomenalism*, which constructs physical objects out of sensory experience alone and says that to see (for instance) a green field is to experience “green-fieldly” in a certain vivid and stable way. To see a thing is to have a visual experience that predictably occurs under certain conditions, say when one has the experiences of walking out on the porch and looking ahead.

On this phenomenalist view, perception can occur without even sense-data; it requires only perceivers and their properties. Sense-datum versions of phenomenalism, however, have been more often discussed by philosophers, and I will concentrate on them.

Whereas the sense-datum theory is an indirect realism, phenomenalism is a *direct irrealism*: it says that perceptual objects are directly perceived, but it denies that they are real in the sense that they are mind-independent and can exist apart from perceivers. This is not to say they are not perceptually real—real items *in sensory experience*. The point is that they are not metaphysically real: things that are “out there,” which are the sorts of things we think of as such that they would exist even if there were no perceivers.

Phenomenalism does not, then, deny that physical objects exist in the sense that they are both stable elements of our experience and governed by

causal laws, such as those of physics. Nor does it deny that there can be hallucinations, as when certain experiences, like those presenting Macbeth's hallucinatory dagger, are too unstable to represent a physical object, or occur in only one mode, such as vision, when they should have tactile elements as well, such as a cool, smooth surface. What phenomenalism denies is that physical objects are real in the classical sense, implying that their existence is independent of our experience.

One naturally wonders why things would not go in and out of existence depending on whether they are experienced, and why, when they do exist, they obey the laws of physics, which certainly do not seem to depend on our minds. Berkeley did not neglect to consider what happens to things when we cease to perceive them, as when we leave a book in an empty room. His answer has been put in the following exchange:

There was a young man who said "God  
Must think it exceedingly odd  
If he finds that this tree  
Continues to be  
When there's no one about in the quad."

Reply:

Dear Sir:  
Your astonishment's odd:  
I am always about in the quad  
And that's why the tree  
Will continue to be,  
Since observed by, Yours faithfully, God.

If the very existence of external objects is sustained by divine perception, it is not difficult to see how their behavior could obey laws of nature that are divinely ordained.

A phenomenalist need not be a theist, however, to offer an account of the stability of external objects and their lawful behavior. John Stuart Mill, writing in the same epistemological tradition as Berkeley but without any appeal to God, considered external objects *permanent possibilities of sensation*. To say that the book is in the room when no one is there to perceive it is to say that there is a certain enduring possibility of the sensations, where having those sensations in a certain stable way constitutes perceiving such a book. If one enters the room and looks in the appropriate direction, that possibility should be realized. By contrast, if one merely hallucinated, there would be no reason to expect this. A phenomenalist can, however, be more radical and take objects not to have any kind of existence when unperceived. They are born and die with the experiences in which they appear.

### *Appraisal of phenomenalism*

Unlike the sense-datum theory of perception, phenomenalism is only occasionally defended by contemporary philosophers. But it has had major influence. Moreover, compared with the sense-datum theory, it is more economical and in that way simpler. Instead of perceivers, sense-data, and external objects, it posits, as the things figuring in perception and sensory experience, just perceivers and sense-data. Indeed, adverbial phenomenalism does not even posit sense-data.

As a theory of perception, then, phenomenalism has fewer objects to analyze and interrelate than do the other theories we have discussed. In addition, it appears to bridge the most important gap between sensory experience and perception of objects: since the objects are internal and directly experienced, it seems natural to say that they must be as they appear to be—we see all there is of the surface facing us and in principle can see all there is to them as physical objects. On the other hand, for the external objects of common sense, whose reality is independent of perceivers, phenomenalism (if nontheistic) must substitute something like permanent possibilities of experience. Thus, the bare-bones appearance of the theory is illusory. Even that metaphor is misleading; for our bodies are also collections of sense-data; even the flesh itself is not too solid to melt into the sensations of its perceivers.

It is tempting to reject phenomenalism as preposterous. But if we do, we learn nothing from it. Let me pose just one objection from which we learn something important about the relation between sense experience and external objects. The theory says that a book, for instance, is—or at least that its presence is (necessarily) equivalent to—one's having or potentially having a suitably stable collection of sense-data, and that seeing it is being visually acquainted with them. If this is a correct analysis of what seeing a book is, then there is a combination of sense-data, sensory items such as colors and shapes in one's visual field, such that if, under appropriate conditions, these elements occur in me, then it follows that I see a book. But surely there is no such combination of sense-data—a point that is important for understanding skepticism. No matter how vividly and stably I (or anyone) may experience the colors and shapes appropriate to a book, it does not follow that anyone sees one. For it is still possible that I am just hallucinating one or seeing something else *as* a book.<sup>10</sup>

This kind of hallucination remains possible even if I have supporting tactual experiences, such as the smooth feel of paper. For even the sense of touch can be hallucinatorily stimulated. Thus, seeing a book is not *just* having appropriate booklike experiences, even if it is *partly* this, and even though, as phenomenalists may hold, there is no experienceable difference between a sufficiently stable combination of booklike sense-data and an independently real material book. Still, if seeing a book is not equivalent to any such collection of sensory experiences, phenomenalism fails to account for perception of ordinary objects. If there are objects for which it holds, they are not the kind we have in mind in seeking an account of perception.

## Perception and the senses

I want to conclude this chapter by indicating some remaining problems about perception. I have already suggested that adverbial theories, sense-datum theories, and the theory of appearing provide plausible accounts of perception, though I consider some version of the first kind *prima facie* best and I leave open that some theory different from all of them may be better than any of them. I have also suggested that some perceptually grounded beliefs fail to be justified, and that, even when justified and true, they can fail to constitute knowledge. There are two further kinds of problems we should explore. One kind concerns observation, the other the relation of perception to the five senses.

### *Indirect seeing and delayed perception*

Observing something in a mirror can count as seeing it. Indeed, it illustrates the sort of thing ordinarily considered seeing something indirectly, as opposed to seeing it by seeing sense-data. We can also speak of seeing through telescopes and other instruments of observation, again indirectly. But what if the object is microscopic and colorless, yet appears to us through our lens as gray? Perhaps we see it, but not quite as it is.

If we see a microscopic object at all, however, there must be some respect in which what we see it by is faithful to it or at least represents it by some relation of causal dependence—sometimes called *functional dependence*. This relation is perhaps more aptly termed a *discriminative dependence*, since perceptual experience seems to vary as a function of certain changes in the object, as where a bird's moving leftward is reflected in a movement of the image in our binoculars, yet in a systematic way that enables us to discriminate it from its environment. But what we see a thing *by*, such as color and shape, need not be faithful in all respects. A green field can look black at night; we are nonetheless seeing it. Moreover, we can see something move in the field even if its color *and* shape are distorted.

How much *correspondence* between an object and our sensory impressions representing it to us is required in order for us to see it (or hear it, touch it, and so on)? There may be no answer to this question that is both precise and highly general. The cases vary greatly, and many must be examined in their own terms.

Observation of faraway objects poses further problems. Consider seeing the nearest star. It is commonly taken to be about four light years away. Presumably we see it (if at all) only as it *was*. For the sense-datum theory, we have a sense-datum produced by it as it was. On the adverbial view, we are sensing "starly" in the way we would if we received the relevant visual stimuli at the time the star produced them. If, however, we see it only as it was, do we see *it*, or just its traces?

Suppose that unbeknownst to us the star exploded two years ago. Is it not odd to say we now see it, as opposed to seeing traces of it (as it was)? The

latter view is preferable, on the ground that if we unqualifiedly see something now, it exists now. But this point is compatible with the view that even though we may see a thing that exists now only *as* it was, we still literally see it now, just as, if we see a cup as elliptical when it is in fact round, we still do see the cup. In the case of the star, *if* the causal connection between it and our sensory experience representing it required no elapse of time, we would discriminate it sufficiently to see it as it now is.

Similar points hold for everyday seeing, since there is still some temporal gap, and for hearing. But if I can see the field only as it was a fraction of a second ago, can I still know that it is now green? I think so, provided there is no reason to believe its color has suddenly changed (but this is something to be reconsidered in the light of our discussion of skepticism in Chapter 13). The same is not clear for the star: may we know by sight alone that it exists now, when it would take about four years for us to realize that the light that had been emitted was its last? This seems doubtful, but it may depend on how likely it is that a star of this kind might have burned out during the period in question. If we knew that such stars last billions of years and that this one is only a few million years old, we might plausibly think we know it still exists. It is plain, however, that understanding perception and perceptual knowledge in these sorts of cases is not easy.

### *Sight and light*

We normally regard seeing as intimately connected with light. But must seeing involve light? Suppose you could step into a pitch-dark room and have precisely the experiences you would have if it were fully lighted. The room would thus *look* to you just as it would if fully lighted, and you could find any unobscured object by looking around for it. Would this not show that you can see in the dark? If so, then the presence of light is not essential to seeing.

However, the case does not establish quite this much. For seeing is a causal relation, and for all I have said you are just vividly hallucinating precisely the right things rather than seeing them. But suppose you are not hallucinating and that if someone covered a coin you see with lead or covered your eyes, you would no longer have a visual experience of a coin. In this case, it could be that somehow the coin affects your eyes through a mechanism other than light transmission, yet requiring an unobstructed path between the object seen and your eyes. *Now* it begins to seem that you are seeing. You are responding visually to stimuli that causally affect your eyes. Yet their doing so does not depend on the presence of light.

### *Vision and the eyes*

In an ominous couplet in Shakespeare's *Othello*, Desdemona's father warns Othello:

Look to her, Moor, if thou hast eyes to see;  
 She has deceived her father and may thee.  
 (Act 1, scene iii)

It would not have occurred to him to question whether there is any way (literally) to see without eyes (figuratively, Othello cannot see well at all, which causes his downfall). But philosophers must sometimes ask whether what seems patently obvious is in fact true. Let us, then, go a step further than treating light as inessential to seeing.

Suppose Emma has lost her eyes in an accident, but a camera is later connected to her brain in the way her eyes were. When she points it in a given direction in good light, she has just the visual sensations, say of color and shape, that she would have had by looking with her eyes. Might this not be seeing? Indeed, do we not think of the camera as *functioning* like the eye? If, under the right causal conditions, she gets the right sorts of sensations through her eyes or a *functional equivalent* of them, she is seeing.

But are even “eyes” (or organs functioning like eyes) necessary for seeing? What if someone who lacks “eyes” could get visual sensations “matching” the objects in the room by strange radiations they emit? Suppose, for instance, that moving the coin away from the person results in the person’s visual impression’s representing a decrease in its size, and that the impressions of it are eliminated entirely by enclosing the coin in cardboard. This confirms the presence of an appropriate causal connection between the coin and the discriminative visual experience of it. If no part of the body (other than the brain) is required for the visual impression of the coin, there is no organ plausibly considered a functional equivalent of eyes, but might we not have seeing?

If what is crucial for seeing an object is its producing visual sensations suitably corresponding to it and appropriately responsive to changes in it, presumably the case is one of seeing. If seeing requires the use of an eye or equivalent organ, then it is not—unless the brain itself is a visual organ. It is clear enough that the person would have knowledge of what we might call visual properties, above all colors and shapes. One might call that visual knowledge. But visual knowledge of this kind could be held not to be grounded in seeing, nor acquired through use of any sense organs. For these reasons, we might doubt whether it must be a kind of *perceptual* knowledge. But a case can surely be made for the visual sensation conception of seeing, as against the organ-of-sight conception.

This case, however, may be challenged: can there be “blind sight,” seeing in the absence of visual sensations? Something like this is reported in the psychological literature. Imagine an ideal case in which a person with excellent blind sight can navigate among obstacles as if the person saw them, while honestly reporting an absence of visual sensations. Could this be seeing?

We automatically tend to understand such behavior in terms of seeing,



and there is thus an inclination to say that this is seeing. The inclination is even stronger if light's reaching the eyes is necessary for the person to avoid the obstacles. But if the subject has no visual sensations—as opposed to lacking ordinary awareness of such sensations—it is not clear that we must say this, and I doubt that it would be true. The most we must say is that the person seems to *know* where the obstacles are. Knowing through some causal process by which objects produce true beliefs about them is not necessarily perception, and certainly need not be seeing.<sup>11</sup>

It may seem that blind sight is genuine seeing because it produces knowledge of “visual propositions”—propositions ascribing visual properties. But knowledge of these propositions is possible without vision, for instance by something like sonar. Moreover, even dependence on light does not establish that the process in question is visual: the light might somehow stimulate non-visual mechanisms that convey information about the objects emitting it. Similar questions arise for the importance of sensations to perception in the other sensory modes, for instance of auditory sensations in hearing. There, too, we find hard questions for which competing answers are plausible.

It is difficult, then, to provide an overall philosophical account of just what seeing, or perception in general, is. All the theories we have discussed can help in answering the questions just posed, but none does so in such a simple and decisive way as to leave all its competitors without some plausibility. Still, in exploring those theories we have seen many important points about perception. It is a kind of causal relation. Even its least complex and apparently most basic mode, simple perceiving, requires, in addition to the perceiver, both an object of perception and a sensory experience that in some way corresponds to that object and records, if only imaginatively, an indefinite and possibly quite extensive amount of information about the object. Partly on the basis of this information, perception tends to produce beliefs about the perceived object. It implies that the perceiver at least normally has justification for certain beliefs about the object, and it normally produces both justified beliefs about that object and knowledge of it.

Perception may be illusory, as when something appears to have a property it does not have, such as ellipticality when it is really circular. Perception—or sensory experience that seems to the subject just like it—may also be hallucinatory, as with Macbeth's dagger. When it is, the question arises whether there must be interior objects, sense-data, with which perceivers are directly acquainted. But both illusions and hallucinations can apparently be accounted for without positing sense-data, and thus without adding a further element to the four that seem central in perception—the perceiver, the object perceived, the sensory experience, and the causal relation between the object and perceiver in virtue of which that experience occurs. Illusion and hallucination can also be accounted for without denying that perceptual experience—the evidence of the senses—normally yields justified belief and knowledge about the world outside the perceiver. Many questions remain, but so far we have

seen no reason to doubt that perception is a rich and basic source of both knowledge and justification.

## Notes

- 1 The theory of appearing has not been widely defended, but a detailed sympathetic treatment is given in William P. Alston's 'Back to the Theory of Appearing', *Philosophical Perspectives* 13 (1999), 181–203.
- 2 That the hallucinatory experience is (in its sensory and other phenomenal properties) qualitatively *exactly* similar to the corresponding veridical one, as a proponent of the argument from hallucination intends, is denied by some philosophers—*disjunctivists*—who hold that there are two qualitatively different elements here (hence a disjunction), not a single experiential element common to both cases. Disjunctivists need not deny the plausibility of the argument, however, but evaluating their case against its first premise would require much discussion. For an informative non-technical treatment of disjunctivism see William Fish, *Philosophy of Perception* (New York: Routledge, 2010), esp. Chapter 6.
- 3 For a recent study and defense of a sense-datum theory see Howard Robinson, *Perception* (London: Routledge, 1994). Cf. Laurence Bonjour's *Epistemology: Classical Problems and Contemporary Responses*, 2nd edn (Lanham, MD: Rowman and Littlefield, 2009).
- 4 The view that ordinary perceptual belief is non-inferential is controversial and—for various senses of 'inference'—has been widely discussed by both philosophers and psychologists. Not *all* sense-datum views, moreover, take perceptual belief to be non-inferential. For a discussion of perception that brings to bear both psychological and philosophical literature see John Heil, *Perception and Cognition* (Berkeley and Los Angeles: University of California Press, 1983), esp. Chapter 2. Cf. Armstrong, *Belief, Truth and Knowledge*.
- 5 Granted, the mouth of the cup does not appear to us *to be* elliptical if we realize its shape cannot be judged from how it visually appears at an angle, but that is a different point. It concerns what shape we *take* it to have, not what shape visually appears in our consciousness antecedently to our taking it to be of any particular kind.
- 6 For a detailed and influential discussion of the adverbial theory, with criticism of the sense-datum view, see R.M. Chisholm, *Perceiving* (Ithaca, NY: Cornell University Press, 1957).
- 7 This is a very important point. One major materialist theory of the mind–body relation—the *identity theory*—says that mental phenomena are identical with brain states or processes. But this theory fails if sense-data exist as mental entities and have properties, such as being green and rectangular, that no brain process has. Identity theorists thus generally oppose the sense-datum theory. See, for example, J.J.C. Smart's much-

discussed 'Sensations and Brain Processes', *Philosophical Review* 68 (1959), 141–56.

- 8 These and other problems are brought against the sense-datum theory by Winston H.F. Barnes in 'The Myth of Sense-Data', *Proceedings of the Aristotelian Society* 45 (1944–5), 89–117. Cf. R.M. Chisholm's discussion of the problem of the speckled hen: is there, for instance, any answer to the question of how many spots there are in an image of such a hen? And how can we distinguish miscounting the number there are from the number's changing as we count or recount? See his *Theory of Knowledge*, 3rd edn (Englewood Cliffs, NJ: Prentice-Hall, 1989), p. 25.
- 9 For a detailed twentieth-century defense of phenomenalism, see Book II of C.I. Lewis's *An Analysis of Knowledge and Valuation* (La Salle, IL: Open Court, 1946). Cf. R.M. Chisholm's widely known criticism of this defense in 'The Problem of Empiricism', *Journal of Philosophy* 45 (1948), 512–17.
- 10 Berkeley might hold that if *God* has booklike sense-data, it does follow that there really is a book. A case can be made for this, but one might also argue that as an all-powerful being God could bring it about that there is a distinction between his creating a physical object and having the corresponding sense-data.
- 11 A subject who really *does* have visual impressions could also misreport. The possibility of such misreporting about one's own consciousness is discussed in Chapter 3.

### 3 Memory

The preservation and reconstruction of the past

- **Memory and the past**
- **The causal basis of memory beliefs**
- **Theories of memory**
  - Three modes of memory*
  - The direct realist view*
  - The representative theory of memory*
  - Memory images*
  - Remembering*
  - The phenomenalist conception of memory*
  - The adverbial conception of memory*
- **Remembering, recalling, and imaging**
- **Remembering, imaging, and recognition**
- **The epistemological centrality of memory**
  - Remembering, knowing, and being justified*
  - Memorial justification and memorial knowledge*
  - Memory as a retentional and generative source*

### 3 Memory

#### The preservation and reconstruction of the past

I believe that I have pruned the flowering crab apple tree which stands in the center of the lawn. This belief is apparently grounded in my memory. When I look at the tree and notice its shape, it often occurs to me that I have pruned it. When this does occur to me, I have a sense of already believing it. The proposition that I have pruned the tree does not seem to be a discovery or a result of inference or a bit of wishful thinking, but rather something I have had in mind before and now believe with some conviction.

On the basis of all these facts about my belief that I have pruned this tree—especially my sense of having already believed this—the belief is justified. Indeed, I cannot help also thinking I know that I have pruned the tree. In particular, my belief that I have seems to be grounded in memory, in the way that what I genuinely *remember* is grounded there. Consider remembering one's having just read the preceding part of this page. If one has just done so, there is likely to be a clear sense of having done it. We do not, for instance, simply have a dreamlike recollection, nor are we concluding what we seem to remember from something else, as you might conclude, from the distinctive shape of a tree, that it must have been you who pruned it.

What, in general terms, is memory? Is it anything beyond a storehouse of some of what we have experienced and learned? And what is it to remember something? Whatever remembering is, its objects include people, material things, facts, events, and, among the events, our own experiences. We might also say that remembering is the chief “function” of memory. Is remembering, then, exercising, or being able to exercise, the capacity of memory? And is there—as with perceptual knowledge by contrast with mere perceptual belief—a special kind of *success* that goes with remembering something as opposed to simply believing it from memory?

In pursuing these questions, it is useful to compare memory with perception. Both are crucial for our knowledge of what is external to the mind: the latter gives us a view of what is outside of us in the present, the former of what is outside of the present altogether. Moreover, memory builds on perception; it preserves much important information we acquire through the senses. It also preserves information about our mental lives. But how does memory achieve this preservation? Must it, for instance, operate by storing

images, or can it preserve bare facts? Before we can see how memory is connected with knowledge and justification, we must first understand what it is and something about how it works.

## Memory and the past

We can learn some basic points about memory and remembering by clearing away some tempting mistakes. To begin with, we cannot say simply that memory is a capacity for knowledge or belief about the past. It is true that memory entails that capacity; but one could have and even exercise the capacity *without* exhibiting memory of the past. Consider the events of World War II. I can know a good deal about them through reading, but at the time I am learning about them through reading I have no memory of them. I witnessed none of them, and I do not remember them. To be sure, I may remember a *description* of them and thereby say—perhaps recalling a history class—that I remember (for instance) the invasion of Normandy. This could be called remembering the events *indirectly*. But it is not remembering in the direct and primary sense that concerns us.

Far from all knowledge of the past being a kind of remembering, then, we commonly know propositions about the past on a basis other than remembering it. Consider again the knowledge of the past obtained while reading; this knowledge is not a case of *remembering* the past but a kind of knowledge of the past acquired through testimony about it. Similarly, I can gain knowledge about the past from your present description of what you did yesterday. This knowledge may not be *retained*, hence need not become memorial. It may never get into the storehouse: I could lose it after I have acquired it, just as we forget a phone number needed only for a moment. In these instances, I have knowledge of the past, but only for too brief a time to qualify as *remembering* the propositions I momentarily knew.

The same example shows a second major point. Like knowledge of the past, beliefs about the past, such as those I acquire about your activities, do not necessarily represent memory. For they need not be *retained* and so are not *memory beliefs*, that is, beliefs grounded in the “faculty” of memory. They are grounded in testimony and are forgotten before being memorially stored.

Moreover, even when one does memorially retain beliefs about the past, they need not amount to *remembering* something. Retained beliefs about the past can be sheer fabrications unconnected with memory capacities. Imagine, for instance, that although I have not seen you for a year, for some reason I groundlessly form the belief that precisely a month ago you wore the belt I see you wearing now. This belief is not memorial, and even when retained, it would not constitute remembering. It comes not from memory of a past event but from undisciplined imagination. Retention of a conviction grounded in fantasy does not upgrade it into remembering.

One might think that beliefs about the past, when they *are* memorial, and not merely retained, constitute remembering. But this need not be so,

because they may be false, whereas everything we genuinely remember to be the case is true. Remembering is, then, *factive*. If, for instance, I remember that Thomas Reid discussed John Locke's ideas about memory, then he in fact did.

To see that even a vivid memory belief that something is so need not represent genuinely remembering that it is, suppose my memory plays a trick on me and I misremember an actual event. I mistakenly think, and vividly imagine, that I planted a green spruce when it was really a blue one I planted. I would now have a memory belief (one tracing back in a normal way to an event it is about) that is mistaken, even though its close relatives in my memory are true. Still, I cannot remember that I planted a green spruce if in fact I did not. Even when a falsehood resides in the neighborhood of truth, it is not elevated into an object of remembering simply because it is retained in a memory belief.

Suppose, however, that a retained belief about the past is true. Is it then an instance of remembering? Not necessarily. Even true beliefs about the past may be utterly baseless and true only by lucky accident. Suppose that my memorially retained belief that you wore that belt just happens to be true because by chance you did select the same belt for both occasions. This belief still does not represent *remembering*. I have merely retained my luckily true impression that you wore the same belt. A retained belief of this sort is stored in memory, but only *properly grounded* true beliefs stored there constitute remembering that something is so.<sup>1</sup>

### The causal basis of memory beliefs

One might think that just as perceptual beliefs are caused by an object perceived, memory beliefs are caused by a past event remembered. Some memory beliefs are caused in this way, and we will soon see that causal connections to the past are essential for genuine remembering. But even if it should be true that all memory beliefs are produced at least partly by events in the past, past events are not the only *objects* of memory or the only things it "stores." We remember, and thereby retain and believe, general truths, such as mathematical theorems. Mathematical propositions are certainly not past events (propositions are not events of any kind). *Learning* them is a past event for many people, but that is a quite different point. Nor are the propositions past objects of some other sort, or even about the past; but many truths of mathematics are clearly among the objects of remembering—the things we remember.

Moreover, even if every memory belief is at least partly caused by a past event, not every belief caused by a past event is memorial. This point applies even if the belief is true. Suppose that my unknowingly taking a poisonous drug causes me to feel strangely ill an hour later, and my feeling ill then causes me to believe I have been poisoned. Then, indirectly, the past event of taking

poison causes me to believe that I have been poisoned. But this belief is not memorial: it is in no way grounded in my capacity for remembering, and I have no memories connected with the belief, such as a memory of someone's putting a white powder into my soup. I arrive at the belief by inference to what I think best explains my illness. Thus, the belief's being caused by the past event of my taking the drug need not make it a memory belief, even if the belief is true. My memory has played no role in supporting the *content* of the belief. The belief lacks a ground appropriate for suitably connecting it with the past event it represents.

An analogy with perception will help. Consider a belief caused by a flash that I do not see, but merely feel as a momentary heat. This belief need not be a visual belief, even if it is a true belief with visual content, say that a camera flashed near my hand. A belief caused by something visible is not thereby a visual belief, just as a belief caused by a past event—something rememberable—is not thereby a memory belief. Since a belief caused by a past event need not be memorial, it is not true that a memory belief simply is a belief at least partly caused by a past event.<sup>2</sup>

The analogy between memory and perception is limited, but it does get us on the right track. For surely a belief about the past is memorial *only* if it has some causal connection to a past event, just as a belief is perceptual (say, visual) only if there is some causal connection between it and the perceived object. Even a belief that arises from testimony and not from first-hand observation and is then stored in memory is *traceable* to the past event of one's acquiring the belief. A thing cannot normally be stored in memory unless it has entered that storehouse. Since memory beliefs can concern any subject, including future events or mathematical truths, we can see that such beliefs need not be *about* any particular event even if their existence does *trace* to one.

Could we, however, have *innate* beliefs? If so, they could be about the past but not memorially connected with a past event, perhaps because the belief is possessed at the time one came into being and does not trace to any remembered experience. It would not *enter* the storehouse of memory: it is part of one's initial equipment. To be sure, perhaps an innate belief could be memorial in roughly the ordinary way if in some *previous* incarnation it traces to an appropriate event, something of the kind that produces a memory.<sup>3</sup> Otherwise, it is best considered merely a *retained belief*, say retained from birth as part of one's native endowment, rather than a belief entering one's memory through, say, observation or testimony.

Just as it is hard to specify how, in order to be perceptual, a belief must be causally connected to the perceptible object it is about, it is hard to specify how, in order to be memorial, a belief must be causally connected to the past. This will become clearer as we explore memory, but fortunately many points can be made about memory without a detailed account of the kind of causal connection in question.



## Theories of memory

If we view theories of memory on the model of three major kinds of theories of perception discussed in Chapter 1, there is much we can discover both about the kinds of causal relations required for remembering and about how memory grounds justification and knowledge. Broadly speaking, the three kinds are direct realism (including the adverbial theory and the theory of appearing as well as naive realism), representative realism, and phenomenalism. Each has an analogue in the theory of memory.

### *Three modes of memory*

In constructing theories of memory, there are at least three different but closely related notions we must track: *memory*, *remembering*, and *recalling*. We have memories of many things. We remember, and we sometimes recall—roughly, call back to mind—much that we have experienced. Both points apply to us in virtue of the power of our memory, conceived as a capacity (a mental “faculty”). There are things we remember, such as isolated facts, that we may never have occasion to recall. But they remain in the storehouse of memory ready to be retrieved if needed. When retrieved, we may be said to have, at the time of retrieval, *a memory* of them.

Our memory, conceived as a “faculty,” is a general capacity: the better it is, the more memories we can have, the better we remember, and the more we can recall. Among the things we remember are *skills* and related behavioral capacities, both mental and physical. Memory of skills is *remembering how*. Much of what emerges here concerning remembering will apply to remembering how (though the notion does not seem reducible to any kind of knowing *that*), but I will not take time to discuss it specifically.

There are, then, three memorial notions to be accounted for by a theory of memory: first, remembering of events, things, and propositions; second, recalling those items; and third, memory as the capacity in virtue of which remembering and recalling occur. There is a further task: accounting for errors. Like perception, memory, as the capacity *for* remembering and recalling—and I include *recollection* as a kind of recalling—can produce impressions that are illusory or, in a way, hallucinatory. Not every memory belief is true; not every recollection is faithful to what it recalls.

In developing the memorial counterparts of the three main kinds of theories of perception, I will concentrate chiefly on remembering, particularly on the simple remembering of events—*event memory*, for instance of my pruning the tree—as opposed to remembering *that* I pruned it, *propositional remembering*, or remembering the pruning *to be* hard, *objectual remembering*. I assume that, like simple perception of something, simple remembering of an event, such as a bird’s flying by, does not entail having a belief about it, as opposed to being disposed to form beliefs about it if the occasion elicits them. But let us concentrate on cases in which one does have such a belief. These cases are crucial for understanding memory knowledge.

### *The direct realist view*

There is a memorial counterpart of naive realism in the theory of perception. It is the view that when we remember an event, we just plain remember it and it is as it seems to us to be. This might be taken to mean that the event is directly presented to us by our memory, as if it were *present* in memory, just as a flash might be present to vision. The difference is that unlike a flash that fills one's visual field, the remembered event is not taken to be occurring. Like all the plausible accounts of memory, this one is best construed as causal: as assuming that some causal chain links us to the remembered event. If I remember seeing Bill a year ago, then it must be in part *because* I did see him that I believe (or am disposed to believe) that I did, and not, say, because I dreamt that I did.

As a direct realist view, this position also maintains that a memory belief is not produced by any intermediary with which we are acquainted, such as an image. To say that would imply a counterpart of the sense-datum theory. We would have an *indirect* realism: just as we perceive the outside world through sense-data that present it to us, we remember the past by virtue of memory's presenting it to us.

Naive realism about memory is inadequate. To begin with, not just any causal connection to the past will do, as we saw with the poisoning case. The causal chain linking a memory belief to a remembered event must be in a sense *unbroken*. In part, the idea is that a belief retained in memory cannot be lost from it during the period of retention. To see the idea, consider a broken chain. Imagine that you saw me prune the apple tree and you remember my doing so. The pruning is then the main causal ground of your memory belief, as it is of mine, and we both remember my pruning it. But suppose I completely forget the event and thus no longer believe I pruned the tree, *then*, solely on the basis of your testimony, later come to believe (once again) that I pruned it. There is still a causal chain from my present belief back to the pruning; for the pruning produced your belief, which in a way produced your testimony, which in turn produced my present belief. But the memorial chain in me was broken by my forgetting.

Given this kind of broken chain to my pruning, I do not retain my original belief and do not remember my pruning; I simply know, from your testimony, that I did it. I now believe the "same thing" but do not have the same (the original) belief.<sup>4</sup> To be sure, *after* your testimony, when I have retained the knowledge you gave me, we might say that I now remember, again, *that* I pruned the tree but no longer remember *pruning it*. Propositional memory *about* an event, even an action of one's own, does not entail event memory *of* it.

The case also shows that knowledge of a past event, even if it is one's own action, does not entail remembering it. I know *that* I pruned the tree, but I do not *remember* pruning it. My propositional knowledge of the event no more represents remembering *it* than my knowledge based solely on your testimony that there is a radiant sunset visible from the front porch represents my

seeing it, when I am inside reading. If, however, I know that I pruned it, then that event is a real element in the past.

A realist view of remembering seems correct, then, if it is coupled with the requirement of an unbroken causal chain. But as stated so far, the view is deficient in some of the ways that naive realism about perception is. For one thing, memory is subject to illusion. I might remember an event, such as meeting you, but not quite as it was, just as you might see white paper in yellow light, and thus not see it as white but as yellow. Here I do not simply remember; I remember incorrectly, for example in remembering the meeting as taking place in New York when it was in fact in Chicago. (I correctly remember meeting you; I do not remember the location of the meeting.) Second, there is the memorial counterpart of hallucination: I may have a vivid image of mailing a letter, and might believe I remember doing so, yet be mistaken. We must, then, account for memorial illusion and similar problems.

### *The representative theory of memory*

The territory may begin to look familiar, particularly if we recall the sense-datum theory of perception, which posits inner sensory objects that, as intermediaries between external things and the mind, represent the former to the latter. For instance, suppose that there are memory images, and that they are genuine objects which figure in remembering rather as sense-data are thought to figure in perceiving. These images might even *be* sense-data if they are vivid enough, but normally they are more like the images of imagination. It might be like this: seeing the apple tree as I prune it produces sensory images in me (whether these are sense-data or not); my memory images of the tree might be conceived as a kind of *residue of perception*.<sup>5</sup>

Perhaps, then, we may be said to remember an event when we have at least one true belief about it suitably grounded in a memory image of it, that is, an image of it which derives, by a suitable unbroken chain, from our experience of the event and represents it correctly in at least some way. The better the memorial representation of the event, the better our memory of it. Call this view *the representative theory of event memory*. It takes event memory to be a representational faculty that works through images that “picture” what they represent.

### *Memory images*

Like the sense-datum theory of perception, the representative theory of memory is an indirect realism. It construes our remembering as *mediated* by memory images (though not as based on inference from facts about such images); it is *through* images that we are acquainted with the past. The view is also like the sense-datum theory in readily accounting for memorial illusion

and similar problems. To remember incorrectly, as opposed to simply having a false belief about the past with no basis in memory, is to be acquainted with a memory image that, despite its being sufficiently faithful to the remembered event to ground one's remembering it, has *some* aspect which produces a false belief about the event, say that it was in New York rather than in Chicago.

The counterpart of hallucination is also treated as one would expect by analogy with the sense-datum theory. Memorial hallucination occurs when one has an image that is intrinsically like a memorial one, but not linked to a past event by a suitable causal chain, just as, in perceptual hallucinations, the sense-data are not produced by the object (or are produced by it in an abnormal way).

Unfortunately, the representative theory of memory has many of the difficulties of the sense-datum theory and some of its own. Consider the similar difficulties first, particularly in relation to remembering.

### *Remembering*

Remembering an event surely does not require acquaintance with an image of it. You may be able to reel off, from memory, some details of a conversation you heard a week ago, even if you have no images, even auditory ones, of the conversation or what it concerned. Moreover, misremembering an event does not require acquaintance with something, such as an image, which actually *has* the property one mistakenly remembers the event as having had, as a sense-datum representing the mouth of a cup from a certain angle is supposed to have the property of ellipticality. I can misremember my meeting you by remembering our meeting as being in New York, when it was actually in Chicago, even if the mistaken element in my memory is not accompanied by images of anything in New York (and even if the correct aspect of my memory is not accompanied by an image that is of our actual meeting in Chicago). I may simply remember the occasion with its animated conversation, yet have the false impression that it was in New York.

Memorial thinking—an episode of thinking about one or more remembered objects or events—may also be possible without objects to serve as images of the past. In retrospective imagination, might I not vividly experience our conversation even if I am acquainted with no object that represents it for me in the way that, in hallucinations, sense-data are supposed to represent physical objects?

Granted, if I have no images, then I cannot *recall*—in the sense of bringing back into my visual consciousness—the color of your sweater. But I might still remember what you said and the hoarseness of your voice owing to the flu, and I might remember what color your sweater was even if I cannot bring the color itself to mind (perhaps you said that its pale blue matched your jacket, and by that remark I remember what the color was without imaging it). I can apparently imagine past events—whether accurately or

inaccurately—without having direct acquaintance with memorial pictures of them, just as I can apparently hallucinate an object without having direct acquaintance with a sense-datum representation of it.

A further difficulty for the representative theory arises when we consider a disanalogy between remembering and perceiving. I can remember our meeting and describe it to someone from memory even if I have *no* images or image-like experiences at all, whereas I apparently cannot see a tree if I have no visual sensations, such as the impressions of foliage that make up an image of a tree. Remembering, even of events that one has perceived, is neither a sensory event nor necessarily an *imaginational* one (even if it often is, especially in some people, such as those who are highly “visual”). So there need not be, in every case of remembering, even the *makings* of a representative theory to which images are crucial.

### *The phenomenalist conception of memory*

The kinds of difficulties we have seen in relation to the representative theory of memory suggest that the memorial counterpart of phenomenism may also suffer irreparable difficulties. Above all, a phenomenalist account of memory relies on images or imaging at least as heavily as does the representative theory, but neither images nor imaging seems either necessary or sufficient for remembering events. Let us explore this.

On what may be the most plausible phenomenalist account of memory, remembering an event is understood in terms of the imaginational content of present experience. To remember an event is (roughly) to have a suitable collection of images representing it, on the basis of which, in a certain way, one believes (or is disposed to believe) something about that event.

This will not do. Remembering an event simply does not require a collection of images analogous to the sense-data from which phenomenists try to construct physical objects (or even a collection of imaging experiences such as an adverbial phenomenist might posit). Images of the kind posited to account for remembering are not only not necessary for remembering, as our examples show; they also are not sufficient for it either. Just as no collection of sense-data is such that its existence implies perception of an external object, no collection of images (even apparently memorial images) is such that, in having a belief about the past grounded on those images, one must be remembering something. No matter how vivid my images of talking with you beneath the skyscrapers of Wall Street, I might not remember our talking there, and my belief that we did talk there (or anywhere) might still be mistaken.

### *The adverbial conception of memory*

If these difficulties are as serious as they seem, then if, in search of a better account of memory, we are to change course and construct a plausible

alternative theory of remembering, we must take account of them. First, such a theory will not claim for remembering all the kinds of directness it posits for perception. Plainly, memory is not *temporally direct*, as past events are not temporally present, whereas we can see a thing's properties at the same time that it has them.<sup>6</sup>

By contrast, any plausible account of remembering, such as a well constructed adverbial theory of it, will take remembering to be (as perceiving apparently is) *epistemically direct*. Memory beliefs, as we have seen, are not inferential. It is not on the basis of any *premise* that I believe (or know) that I have pruned that crab apple tree. My belief is grounded in memory as a preserver of beliefs and other elements, not in other beliefs which give me premises to support the memory belief.<sup>7</sup>

Moreover, an adequate theory must not say that (actively) remembering an event, such as pruning a tree, is constituted by memorially imaging in a way suitably caused by that past event, as perceiving an object is sensory stimulation suitably caused by the thing perceived. For no such imaging need occur (though it commonly does, especially in the case of the active remembering that constitutes recalling). We can describe a past event to others, and in doing so actively remember it, even if we are imaging nothing but their faces.

Positively, the *adverbial view of memory*, applied to remembering events, should be expressed as something like this. First, *actively (occurrently) remembering* an event is *manifesting—realizing*, in a sense—a memorial capacity concerning it, in which this manifestation is linked to the event by an unbroken causal chain. Just as, in observing a cat, one is manifesting a perceptual capacity, in describing a play from memory one is realizing a memorial one. The most typical manifestations—the things that constitute experiencing in a memorial way—are probably (1) imaging processes concerning the event; (2) formations of memory beliefs about it, often through considering one's memory images; and (3) considering the propositions so believed, with a sense of already believing them. But there may be other realizations, for instance recognizing a picture of the event. Second, *passively (dispositionally) remembering* an event is having this capacity in an unmanifested way, as when, though I can recall the pruning if I want to, my mind is wholly on other things. For me to remember the pruning actively, something must call it to mind.

To see the difference between the dispositional and the occurrent in another context, consider elasticity in a rubber band. It is a dispositional property, whereas stretching is an occurrent property that manifests the disposition (this distinction is further discussed in Chapter 4). Recalling an event can be related to dispositionally remembering it much as a thing's stretching is to its elasticity. Just as stretching manifests the disposition of elasticity, recalling is a case of actively remembering that manifests the dispositional memory which retains the thing recalled.

Propositional remembering—remembering *that*—can be construed similarly. On the adverbial view imagined, to remember that an event occurred

is a memorial *way* of truly believing that it did, roughly, to have one or more true beliefs about it which are suitably linked by an unbroken chain to past experience and represent the event in a certain way (if only as occurring). These beliefs constitute knowledge that is preserved in memory. On this view, then, remembering that something is so constitutes knowledge *from memory*, rather as seeing that the cat is sleeping constitutes knowledge through perception.

Most of what we propositionally remember is dispositional, roughly, recorded in dispositional memory beliefs. When these beliefs are called up in active propositional remembering, as when I describe how I pruned the tree, one is experiencing in a memorial *way*. This does not require being acquainted with imagistic memorial objects. One may, but need not, image memorially, as when one actually calls up the remembered experience and focuses on its features in one's imagination.

Moreover, whether one images a remembered event or not, the event need not be entirely *as* one remembers it. Here event memory differs from propositional memory; the former, like seeing *as*, can misrepresent the thing in question, whereas remembering *that* something is so, like seeing that it is, entails its being (exactly) so. One can remember a meeting as being in the wrong city, thus remember it in the wrong way geographically, just as one can see a circular cup as elliptical and so see it in the wrong way visually.

In neither case of illusion, to be sure, does one have to be fooled: with memory as with perception, illusion does not always produce false belief. Typically, if I remember something as having a certain quality, say a conversation as being rushed, I believe it was like that; but I can remember it *as* such, yet know from independent evidence (such as testimony) that it was not rushed. If, however, we really remember some object or event, then we are right about some aspect of it, or are at least in a position to form some true beliefs about it on considering the matter. This is parallel to the point that if one really sees something, one is at least in a position to see it to *be* something or other.

## Remembering, recalling, and imaging

So far, the adverbial view seems superior to its competitors in relation to the crucial notions to be accounted for, such as remembering and recalling. Will this direct realist view stand scrutiny? In answering this, it is important to see that the view can account for imaging; it simply does not take imaging to be an acquaintance with inner objects. Still, there may be a nagging doubt about whether it does not incline us to posit too little imaging. When I am remembering an event, especially a perceived one like a ship's docking as opposed to an imperceptible one like thinking about knowledge, I typically *do* image some aspect of it. I refer, of course, to active remembering, as opposed to my stored remembering of events that are now far from my mind but which I *could* actively recall if the subject came up. The first kind of



remembering is occurrent, as it is in part a matter of something's occurring in me. The second kind is dispositional, as it is a matter of my being disposed (roughly, tending) to remember a thing actively (occurrently) *provided that* something activates my memory. Thus, although yesterday's concert may be far from my mind while I write a letter, if someone asks how I liked the Chopin, then my dispositional memory may be activated; and, as I recall it, thereby occurrently remembering it, I may say I thought it inspiring.

It is occurrent remembering that is analogous to perception and is of chief concern now; and it is occurrent remembering that is closely associated with imaging. Does occurrent remembering require some sort of imagery after all, even if not images as sense-datum objects?

Here is a natural way to answer. Consider one of your memories of an event, for instance meeting someone for the first time. Do this in such a way that you take yourself to be actively remembering that event. Second, ask yourself whether you are now imaging. When I do this, I image. Here, remembering involves imaging. But notice what has happened: I have *called up* a memory and inspected the results of my effort. Perhaps I am imaging because of the *way* I evoked the remembering, or because I scrutinized the process of my calling up the meeting. Self-conscious evocation of the past and scrutiny of the results may yield findings unrepresentative of remembering in general.

This procedure of evoking memories of the past, then—selecting them by recalling past phenomena—is defective as a way of determining whether remembering requires imaging. But the procedure does show something. For suppose that what I have done is to *recall* a past event. Perhaps recalling, which is calling back to mind, often by a lengthy search of one's memory, *does* require imaging provided it is a recalling *of* an imageable event, such as pruning a tree, as opposed to, say, recalling a theorem. There is some reason to think this. If no imaging of our seaside luncheon enters my consciousness, how can I have recalled it? Sometimes, moreover, we say that we cannot recall someone, meaning not that we do not *know* who the person is, but that we cannot image the person. There, recalling seems to imply imaging. To be sure, recalling the content of our conversation is possible simply by becoming conscious, in a certain recollective way, of its informational content. But the content of a conversation, as distinct from our activity of conversing, is not imageable in the first place.

Even if the most common kinds of recalling should imply imaging, however, remembering does not. Why, then, does that idea persist? For one thing, when we collect specimen memories in order to examine remembering, we often do it *by* recalling things. If so, it should be no surprise that the specimen memories involve recalling something. If, in trying to determine the shades of beech leaves, I collect specimens only from the nearby copper beeches, it is no surprise that I may erroneously think beeches in general have copper-colored leaves.



## Remembering, imaging, and recognition

A related point is that what we cannot recall we often *believe* we cannot remember. On the adverbial view, this is natural; for an inability to remember is a lack of a memorial capacity, and, understandably, we may think we lack that capacity when, under normal conditions, we cannot exercise it in an expectable way—such as recalling an event we have been taking ourselves to remember. But imaging is only *one* exercise of memorial capacity, important though it is; and just as we can be capable of climbing a mountain, but not necessarily by every route to the top, we may have the capacity constituted by remembering something, but be unable to exercise the capacity in every way it can be exercised. Hence, inability to image does not imply that one does not remember the thing in question. We can see, then, both why there is a tendency to think that remembering requires imaging and why we should not accept this view.

Imaging may still seem more important for remembering than so far granted. But take another case. Suppose I can neither recall nor image Jane. I might still remember her; for on seeing her, I might *recognize* her, and might even recall, our last meeting. This would suggest that my memory simply needed to be “jogged.” In adverbial terms, before I see her again I dispositionally know her in a certain memorial way—I remember her—even though I cannot imagistically experience her in a memorial way—namely, recall her.

I choose the example of remembering a person because it is easy to show that we do remember someone by creating the right occasion. Recalling her is an indication of my remembering her, but it may not be possible despite my remembering her; recognizing her when I meet her is a proof of the pudding. We cannot draw this contrast with past *events*, since, unlike people, they cannot be literally brought back. But even here, there is indirect recognition, as when one recognizes a ship’s docking in Helsinki harbor upon seeing a picture of the event. It is doubtful, then, that the relation between recalling and remembering is different with events.

It is important to see that the way I am now considering the relation between recalling and remembering is direct and non-introspective. I am exploring what is possible and what it would show. It is possible, however unlikely, that I might have no retained image of pruning the apple tree, yet be able to give an account of the pruning that is both remarkably accurate *and* grounded by a suitable causal chain in the original experience of the pruning. If I do this without my having received any information about the event from anyone else, it is an excellent reason to think I remember the event. It is akin to recognition of a person one could not recall, say by picking the person out of a crowd.

To be sure, our *beliefs* about what events we remember may depend on what we can recall, which, in turn, may be largely dependent on what we can image. But what events we *do* remember is a matter of how our memorial capacities are grounded in the past and not of what kind of evidence we can get, imagistically or otherwise, concerning that grounding.

In exercising my capacity to remember events, then, I need not rely on my images or even on my ability to image, though in fact retention of images doubtless aids remembering. The representative theory of memory therefore seems mistaken, and some memorial analogue of direct realism regarding perception is apparently preferable. The possibility of a good analogy is already implicit in the point that an event can be perceived even though the time at which one has a sensory experience representing it is later than its occurrence. The suggested adverbial view of remembering is a good position from which to work; but I leave some important questions about memory unexplored, and it would be premature to present that view as clearly correct.

### **The epistemological centrality of memory**

We can now see some points about memory as a source of belief, knowledge, and justification. Let us start with beliefs. Memory is a source of beliefs in the way a storehouse is a source of what has been put there, but it is not a source of beliefs in the generative way perception is. Clearly our memory, as a mental capacity, is a source of beliefs in the sense that it *preserves* them and enables us to *call them up*. It also enables us to *draw on our beliefs to supply premises* in reasoning. We do this when we solve mathematical problems using memorized theorems. We may also be guided by other kinds of presupposed premises without having to call them to mind. Remembered propositions (and patterns) can be like routes we know well: they can guide our journey while we concentrate on the road just ahead.

When our memory beliefs are of propositions we remember to be true, they constitute knowledge. If you remember that we met, you know that we did. Similarly, if you remember me, you know me, at least as I was once—you may not be able to recognize me now. So memory, when it is a source of what is remembered, commonly yields both knowledge *that* and knowledge *of*. Remembering, then, is *knowledge-entailing*. The analogy to perception is significant here too.

Is memory also a source of justification? Surely what justifies the great majority of my justified beliefs about the past is my memory. For instance, my belief that I twice pruned the crab apple tree is justified because of the way that belief resides in my memory. It has, for example, a special kind of familiarity, confidence, and connection with other things I seem to remember. Moreover, it appears that if I remember that I met you, I am justified in believing I met you. It thus seems that when memory yields genuine remembering it also yields justification. Certainly this commonly holds.

### ***Remembering, knowing, and being justified***

Perhaps, however, I could remember that I met you, yet *fail* to be justified in my belief because (in fun) you convince me, by good arguments and by enlisting the corroboration of plausible cohorts, that I am probably confusing you

with someone else. Still, suppose my belief remains properly grounded in my actual memory of having met you, perhaps because the memory is so clear that the belief is almost unshakable. Then I may still genuinely remember that I met you. Despite this point, if your arguments are good enough, I may properly reproach myself for still holding the belief that I met you, and my belief may perhaps cease to be justified. Its justification would be defeated by your arguments and by my own credible self-reproaches based on seeing their plausibility.

If this case is possible, it has an important implication. If, as I have suggested, remembering that something is so entails knowing it is so, then the case as described implies that knowing that something is so does *not* imply justifiedly believing it. (In Chapters 11 and 12, I return to the relation between knowledge and justification, but it is important here to see that the domain of memory provides a challenge to understanding that relation.)<sup>8</sup>

Furthermore, if the case is possible and one *can* remember that something is so, yet fail to be justified in one's believing that it is so, then we might question whether memory yields any justified beliefs after all. Fortunately, the example by no means rules this out. Quite apart from cases of genuine remembering, memory often yields justified belief. If I have a vivid and confident belief that I met Jane, and this belief seems to me to arise from a memory of the occasion, I may, simply on that basis, be justified in the belief. Surely this is, after all, just the sort of belief that usually does represent remembering; in any case, I have no reason to question its credentials.

Memory can justify a belief even when that belief does not constitute knowledge or rest on actual remembering of the proposition or event in question. If, for instance, I do not in fact remember meeting Jane, perhaps the only reason why I do not is that it was her identical twin, of whose existence I had no idea, whom I met. That excusable ignorance may prevent my *knowing* that I met Jane, but it does not preclude my justifiedly believing that I did.

### *Memorial justification and memorial knowledge*

These reflections suggest a *memorial justification principle for events: normally, if one has a clear and confident memory belief that one experienced a given thing, then the belief is justified*. We might also call clear, confident memory beliefs *prima facie* justified.<sup>9</sup> A memory belief is one grounded in memory; this is commonly a kind of belief which represents the event or proposition in question as familiar in a certain way. Commonly, if one considered the matter, the belief would seem to one to arise from one's memory; but the notion of a memory belief cannot be defined by that normal property of such beliefs, and it is not easily defined at all.<sup>10</sup> We can believe—even know—from memory propositions we do not find familiar (as when we have not recalled them or thought of them in years).

A still broader principle may perhaps be true—the *general memorial justification principle: normally, clear, confident memory beliefs with any subject*

*matter are prima facie justified.* Moreover, if they do not conflict with other beliefs one holds, say that one has never been to the country where one now seems to remember going to a museum, they tend to be justified on balance. With both principles the degree of justification may not be great, particularly if there is no corroboration, such as apparently recalling a sequence of events related to the belief. My belief that I met someone at a restaurant tends to be better justified if I apparently remember related events, such as a friend's recently mentioning our meeting that person there, than if it is isolated from other apparent memory beliefs that confirm it.

Both these and similar principles help to describe how memory is plausibly conceived as a source of justification. This is certainly how it is standardly conceived. Imagine someone saying "I have a clear and confident memory belief that we met at the Café Rouge, but this gives me no justification whatever for thinking so." We can understand someone's holding that there is *better* justification for not believing this—say, because of known memory failure—but that would show only that the justification is defeated, not that there is none whatever to be defeated.

### *Memory as a retentive and generative source*

There is a very important difference between the way in which memory is a source of knowledge and the way in which it is a source of justification. To see this, we must take account of several points. Memory is a *preservative* capacity with respect to both belief and knowledge. First, when you initially come to believe something, you do not (yet) remember it. Second, you cannot (propositionally) remember something unless you *previously* knew or at least believed it, for instance perceptually, and your belief of it is suitably preserved.

Thus, memory *retains* belief and knowledge. Retention is roughly equivalent to preservation but has a lesser implication of unchangingness; a belief held with considerably less confidence than originally, for example, is less properly said to be preserved than to be retained. Memory does not generate belief and knowledge, except in the sense that, by *using* what you have in memory, you can acquire beliefs and knowledge through inference (or perhaps through other processes that themselves yield belief and knowledge). I may, for instance, infer much from propositions I remember, or I may arrive at greater knowledge of a movie I saw by calling up images of various scenes. Here it is thought processes—inferential and recollective—that, *partly* on the basis of retained material, produce belief and knowledge.

To say that memory is not a generative source of knowledge is not to deny that memory is sufficiently connected with knowledge to figure in a plausible epistemic principle—call it the *memorial knowledge principle*: *normally, a true memory belief, supported by a vivid, steady experience of recall that is in turn corroborated by other memory experiences, represents knowledge.* But if this principle is correct, that is *because* such beliefs are of a kind that ordinarily

constitutes knowledge originally, say when one learned through perception the truth that a tree has recently been pruned, and continued therefore to have grounds, preserved by one's memory, for holding this belief.

Memory is not, then, a *basic source* of belief or knowledge, a source that generates them other than through dependence on a contribution by some different source of them. It is, however, a basic source of justification. We can be justified in believing something either on the basis of remembering that it is so, or on the basis of our having a clear and confident memory belief that it is, even if this belief is not true. If, however, we genuinely remember that something is so, it is so, and we know that it is.

Memory can not only generate justification, as when my vivid sense of remembering a line of poetry justifies believing that I do. Memory can also preserve justification, particularly when the justification resides in memorially retained beliefs, for instance beliefs of one's premises for a view one holds. But the original justification of a belief *need* not be retained in order for the belief to be memorially justified. The sense of memory can generate justification by virtue of the way the proposition or event in question occurs to us, and this could occur when we have forgotten our original, remembered justification, such as witnessing the event that, perhaps because we have read a vivid account of it, we now have a memorial sense of remembering.<sup>11</sup>

This justifying capacity of memory often operates even when we have no associated images. But in accounting for what justifies memory beliefs, images do have a significant if restricted role. We are better justified in a memory belief supported by imagery, especially vivid imagery, than in memory beliefs not thus supported (other things being equal). Perhaps the reason is that we have at least some justification for believing that there is less likelihood of error if both imagery and beliefs point in the same direction, say to my having met you two years ago.<sup>12</sup>

For all the analogy between memory and perception, then, there are important differences. Although both are essential to our justification for believing a huge proportion of what we believe, perception is more fundamental in a way that is crucial to the development of our outlook on the world. It supplies memory with much of its raw material, whereas memory, though it guides us in seeking what to observe and, in that way, often determines *what* we perceive, does not supply raw materials to perception: it manufactures no perceptibles. It does, to be sure, supply raw materials for introspection and thought: we would have vastly less to "look in on" or think about if we did not remember sights and sounds, conversations and embraces, ideas and plans.

Both memory and perception, however, are to be causally conceived, and both are, in different ways, sources of belief, justification, and knowledge, propositional as well as objectual. But perception is a basic source of all three: it can produce them without dependence on contributions from another belief-producing capacity, such as reasoning. Memory, being a capacity for

the preservation, and not the creation, of belief and knowledge, is not a basic source of them.

Still, without memory, perceptual knowledge could not be amassed and used to help us build theories of the world or of human experience, or even to make local maps to guide daily living. We would not even have a sense of who we are, since each moment in our lives would be dead to us by the next. Beyond this, memory is a basic source of justification. That is a vitally important epistemological point. And as we shall see, the role of memory in our knowledge in general is also of enormous epistemological importance.

## Notes

- 1 We might call merely retained beliefs *weakly* grounded in memory, but I reserve the terms ‘memory belief’ and ‘memorial belief’ for beliefs grounded in the normal way illustrated by remembering what I come to believe from, say, perception or testimony.
- 2 The point that how beliefs are caused, and what their content is, may not indicate how they are grounded (where grounding is the notion crucial to determining whether the belief is justified or represents knowledge) is even wider than so far suggested. A noise too faint for me to hear may cause Tom to jump, which in turn causes me to believe that he is startled; my belief that he is startled is thus (indirectly) caused by the noise, but it is not auditory. It is in no way grounded in my hearing.
- 3 In both Western philosophy—for example in Plato and Descartes—and Eastern philosophy, innate ideas have played a significant part. In recent times there has been much skepticism about whether they—as opposed to innate dispositions to *form* ideas—are even possible. I cannot discuss this issue here but see no reason not to leave the matter open for the sake of argument. In any case, the possibility of “innate” beliefs seems implicit in something less controversial: that in principle a person could be created as a perfect copy of another, and so would have at least some beliefs at the moment of “birth.”
- 4 The belief is the same in the sense of being an instance of the same cognitive property—the one we both instantiate—as is indicated by our each believing the proposition; but the *instance* of that property is not the same, much as if I raise my hand twice in the same way there are two (act) tokens of the same (act) type.
- 5 John Locke, in (e.g.) his *Essay Concerning Human Understanding*, speaks of perception as “the inlet of all the materials of” knowledge (Book II, Chapter IX, section 15) and says, comparing perception and memory, that:

when my eyes are shut or windows fast, I can at pleasure recall to my mind the ideas of light, or the sun, which *former sensation had lodged in my memory* . . . there is a manifest difference between

the ideas laid up in my memory . . . and those [of perception] which force themselves upon me . . . there is nobody who doth not perceive the difference in himself between contemplating the sun, as he hath the idea of it in his memory, and actually looking upon it: Of which two, his perception is so distinct, that few of his ideas are more distinguishable from one another.

(Book IV, Chapter XI, section 5; italics added)

- 6 Or virtually the same time: the time-lag argument discussed in Chapter 2 indicates that if light transmission is essential to seeing, there will be a tiny gap between (1) the time at which something we see has a property we are visually caused to believe it has and (2) the time at which we see it as having, or believe it to have, that property. We also found, however, that light transmission does not seem *absolutely* essential for seeing. Note that any causal theory will imply a time-lag when the needed causal connection requires time; but, as with an engine pulling a train, the production of effects by causes need not in all cases require temporal passage.
- 7 I assume that simple inferences do not require the use of memory; but even if they do, once a belief is *formed* inferentially, it can be inferentially *held* only insofar as it remains supported by the premise beliefs. Then memory may well be what preserves the *inferential structure* represented by believing *p* on the basis of premises; but the belief that *p* is itself only preserved by memory without being genuinely memorial. Not every way that memory preserves a belief renders the belief memorial, and one would explain why one holds this belief not by saying 'I remember . . .' but by citing one's premises.
- 8 I develop this case, defend the conclusion tentatively stated here, and discuss other matters considered in this chapter in 'Memorial Justification', *Philosophical Topics* 23, 1 (1995), 31–45. For a different position on some of the relevant issues see Carl Ginet, *Knowledge, Perception, and Memory* (Dordrecht: D. Reidel, 1975).
- 9 It is natural to wonder whether the degree of justification normally belonging to such memory beliefs is as great as that normally belonging to perceptual beliefs. Perhaps not, and one could add 'to some degree' in the normality formulation. But it still appears that the kind of justification is such that it is generally reasonable to believe the propositions in question and that when they are true we commonly can know them on the basis of the relevant kind of justifier.
- 10 My 'Memorial Justification' (cited in note 8) and some of the literature it refers to consider this difficult question; fortunately, it is not one that requires here any more than the sketch of an answer given.
- 11 Cf. Michael Huemer's conception of memorial justification in his 'The Problem of Memory Knowledge', *Pacific Philosophical Quarterly* 80

(1999), 346–57 reprinted in his *Epistemology: Contemporary Readings* (London: Routledge, 2002).

- 12 Note that ascribing justificatory power to memorial images and other memorial experiences, such as the sense of having believed a proposition, *p*, does not commit one to holding that each *time* we have the relevant experience we get *more* justification for *p*, so that we could enhance our justification simply by repeatedly calling up confirmatory images. Granted, a more vivid recollection of the event that *p* represents, such as pruning a tree, may, other things being equal, better justify believing one pruned it than does a less vivid image. We should distinguish *occurrent justification*—the kind based on an experience one is now having—from *dispositional justification*—the kind one has for a belief simply retained in memory. The former varies in degree with the justificatory force of the ground in one's experience at the time; the latter either is invariant, as when one's ground is of a constant justificatory power, or changes only as one acquires new grounds or loses one or more actually possessed. New grounds include new confirmatory evidences; justificatory losses include forgetting or having the confirmatory images one can recall become less vivid. Variability in the force of occurrently justificatory elements does not entail an additive effect. For instance, regarding my belief that I pruned the spruce, calling up my image of pruning it may provide more support on one occasion than on another, perhaps because of differing levels of concentration; but this does not entail that I can parlay my recollections into certainty by simply repeating the exercise, any more than verifying that a pillar remains firmly attached to the corner of one's porch strengthens its support. Indeed, repeating the exercise of recalling an event *can* tend to reduce our justification, at least when it amounts to inflating our self-assurance. I thank Sven Bernecker for raising the difficult question of the justificatory force of repeated recollection.





## 4 Consciousness

### The life of the mind

- **Two basic kinds of mental properties**
- **Introspection and inward vision**
- **Some theories of introspective consciousness**
  - Realism about the objects of introspection*
  - An adverbial view of introspected objects*
  - The analogy between introspection and ordinary perception*
  - Introspective beliefs, beliefs about introspectables, and fallibility*
- **Consciousness and privileged access**
  - Infallibility, omniscience, and privileged access*
  - Difficulties for the thesis of privileged access*
  - The possibility of scientific grounds for rejecting privileged access*
- **Introspective consciousness as a source of justification and knowledge**
  - The range of introspective knowledge and justification*
  - The defeasibility of introspective justification*
  - Consciousness as a basic source*

## 4 Consciousness

### The life of the mind

So far, we have talked mainly about beliefs concerning things outside ourselves: the green of grass, the smell of roses, the feel of glasses in our hands. But there is much that we believe about what is internal to us. I believe that I am *thinking* about self-knowledge, that I am *imaging* cool blue waters, and that I *believe* I am a conscientious citizen. In holding these beliefs, I attribute rather different sorts of properties to myself: thinking, imaging, and believing. What sorts of properties—or at least phenomena—are they, and how do our beliefs about them give us justification and knowledge? For instance, are some of these self-directed beliefs the products of a kind of *inner perception*? This seems a natural view, and we have already seen how an understanding of perception can clarify memory. If there is some truth in the inner perception view of self-knowledge, exploring the analogy between outer perception and self-consciousness might help to explain how beliefs about our inner lives are justified or constitute knowledge.

Our most important kind of self-knowledge is not about our bodies, but about our minds—for instance about what we believe, want, feel, and take ourselves to remember. It will help to start by describing the kinds of mental properties illustrated by thinking, imaging, and believing. Since they are all broadly mental, this is a task in the philosophy of mind. But epistemology cannot proceed without considerable reflection on mental phenomena, and here it overlaps the philosophy of mind. Thinking, inferring, and believing, for example, are central in both branches of philosophy; and to understand self-knowledge, we need a good sense of what kinds of properties characterize us. We might begin with two kinds that, for our purposes, yield a basic division.

#### Two basic kinds of mental properties

Thinking is a kind of *process* and involves a sequence of events, events naturally said to be in the mind. Thinking in human beings has a beginning, a middle, and an end; it is constituted by mental events, such as considering a proposition; and these events are always ordered in time, often in subject matter, and sometimes in logic.

Simply *having* an image, in the minimal way one does when there is a static, changeless picture in the mind's eye, is being in a certain (mental) *state*. Unlike something that changes, the existence of such a state does not absolutely require the occurrence of any events during the time it exists. *Imaging* can be a process of calling up a succession of images or, as when one of them is held changeless in the imagination, static. I could image something for a time without any change whatever in my imaging, and without the occurrence of any mental event that might be part of the imaging.

Believing could also be called a mental state, but this terminology can be misleading in suggesting that having a belief is a state of mind, where that implies a global mental condition like worry or excitement. Unlike images and aroused emotions such as jubilation, beliefs do not tend to crowd one another out.

Beliefs differ from images in at least two further ways. First, beliefs need not be in consciousness and indeed we can be conscious of only a quite limited number at once. We all have many which, unlike my belief that I am now writing, we cannot call to mind without some effort. Second, believing need not in any sense be "pictorial." Consider a belief present in consciousness, in the way my belief that the rain has stopped is. This belief is present because I have called it to my attention. I might have held it without attending to it or even to the fact it records.

Even a belief present in consciousness in a prominent way *and* about something as readily picturable as the Statue of Liberty need not involve anything pictorial in the way my imaging must. Suppose I believe that the Statue of Liberty has a majestic beauty standing high in the Bay of New York. Without picturing anything, I can entertain this proposition, and in that way have this belief in my consciousness. By contrast, imaging cool blue waters requires picturing a blue surface. To be sure, when we *call up* this belief about the statue, we tend to picture that structure. But I could later get the proposition in mind, as when I am listing some majestically beautiful landmarks deserving preservation, without picturing anything. I could even retain the belief if I had forgotten what the statue looked like and simply remembered my aesthetic judgment of it.

It will help in sorting things out if we observe a distinction that has already come up but needs more development. Let us call mental properties like beliefs *dispositional* and mental properties like thinking (processes-properties, we might say) *occurrent*. The latter are constituted by mental events and are occurrences: they take place in the way events do and may be said to happen or to go on. The former are not occurrences and may not be said to happen, take place, or go on.<sup>1</sup>

The basic contrast is this. To have a *dispositional property*, or (perhaps not quite equivalently) to be in a dispositional state, is to be disposed—roughly, to tend—to do or undergo something under certain conditions, but not necessarily to be actually doing or undergoing or experiencing something or changing in any way. Thus, my believing that I am a conscientious citizen

is, in part, my being disposed to say that I am one, under conditions that *elicit* that sort of verbal manifestation of my belief, such as your asking me whether I intend to vote. Yet I can have this belief without doing or undergoing anything connected with it, just as sugar can be soluble while it is still in a solid, unaltered lump. I can have the belief in dreamless sleep. By contrast, to have an *occurrent property* is to be doing, undergoing, or experiencing something, as sugar undergoes dissolution. Thus, if you are thinking about mental phenomena you are doing something, even if you are in an armchair; and if you are imagining a flowering crab apple tree, you are experiencing something, at least in the sense that your imagining the tree is now present in your consciousness.

Having a static image, however, as opposed to *calling up* an image, is not a process as, say, silently talking to oneself is. Occurrent mental properties, then, must be subdivided. To mark the difference, we might call occurrent mental properties like thinking *experiential process properties* and occurrent mental properties like having a static image in mind *experiential state properties*.<sup>2</sup> Clearly, both differ from dispositional mental properties; possessing those does not even require being conscious, much less having a kind of experience. All three kinds of mental properties turn out to be important for understanding the epistemological role of introspection.<sup>3</sup>

## Introspection and inward vision

If we take a cue from the etymology of ‘introspection’, which derives from the Latin *introspicere*, meaning ‘to look within’, we might construe introspection as attending to one’s own consciousness and, when one’s mind is not blank, thereby achieving a kind of inner seeing. I might introspect my images, for instance, and conclude that my image of the spruce and nearby maple indicates that the spruce is taller than the maple. I might *have* to introspect my image of the maple to tell without looking back at it whether it has three secondary trunks. Introspection need not, however, be labored or even constitute an act. It may be simply a matter of becoming conscious of something in my mind. This can be as natural as something’s coming into one’s physical field of vision, rather than like making the effort of observation in order to see.

It is not only in consciously introspecting that one can vividly image. In Shakespeare’s *King Lear* there is a scene in which Edgar wants to convince Gloucester, who has lost his sight, that he is at the top of a cliff. Edgar’s description is so vivid that the deception succeeds:

How fearful and dizzy ’tis to cast one’s eye so low!  
 The crows and choughs that wing the midway air  
 Show scarce so gross as beetles. Halfway down  
 Hangs one that gathers samphire, dreadful trade!  
 Methinks he seems no bigger than his head.

The fishermen that walk upon the beach  
 Appear like mice, and yond tall anchoring bark . . .  
 Almost too small for sight.

(Act IV, scene vi)

What Gloucester sees in his mind's eye is so vivid that he believes he is at the edge of a precipice. His visual consciousness is filled with images from Edgar's portrait. Here introspection is simply a matter of vivid consciousness of the imagery that is before the mind.

If introspective consciousness does produce inner seeing and other sensuous imagery (such as, commonly, sound), we can try to understand it by drawing on what we know about perception. For instance, we can explore introspectional counterparts of some theories of perception and sensory experience. But one limitation of that procedure is apparent the moment we reflect on the dispositional mental properties, for instance believing, wanting, or having a fear of cancer. We do not *see* such properties in any sensory way, nor even as we may be thought to see (in our mind's eye) an image of cool blue waters.

The analogy to vision might, however, still hold for introspection regarding *occurrent* mental properties. If it does, it presumably applies only to the mental state properties, such as imaging. For surely thinking is not seen. It need not even be heard in the mind's ear. I may hear my silent recitation of Shelley's 'Ozymandias', but thinking *need* not occur in inner speech, certainly not speech of that narrative, punctuated sort.

Perhaps it is only pictorial mental properties that we see through inner vision; and perhaps it is only other "phenomenal" properties, such as inner recitations, tactual imagings (say, of the coldness of a glass), and the like, that seem accessible to inner analogues of perception: to hearing in the mind's ear, touching in the tactual imagination, and so on. It is doubtful, then, that we can go very far by conceiving introspection as simply producing inward seeing. Still, it is worth exploring how the analogy to seeing holds up for the one important case of pictorial properties.

## Some theories of introspective consciousness

Suppose that introspecting such things as images of cool blue waters does produce a kind of inner seeing. Are we to understand this seeing on realist lines, so that there must be some real object seen by the introspective eye?

### *Realism about the objects of introspection*

One might think that the sense-datum view simply cannot be extended in this way to introspection. This is at least a natural assumption about self-understanding. For on the introspectional counterpart of the sense-datum view, seeing (in one's mind's eye) an image of cool blue waters would require

something like *another* image, one that represents the first one in the way sense-data represent a physical object seen by virtue of the perceiver's acquaintance with them. Call it a *second-order image*, as it is an image of an image.

What would second-order images be like? If I try to image my image of cool blue waters, I get that very image all over again, or I image something else, or I get something that is not an image at all, such as a *thought* of my original image. But this point does not show that there could not be second-order images. Perhaps there could be some that are less vivid than the originals they picture, just as my imaginational image of blue waters is less vivid than the sensory image I have in seeing those waters. In any case, sense-datum theorists could avoid positing second-order images. It would be more plausible to hold that to image simply *is* to be acquainted with sense-data of a kind different and typically less vivid than outer perception yields.

### *An adverbial view of introspected objects*

A defender of an adverbial account of sensory experience would not countenance images as sense-datum-like entities with properties in their own right. Take first perceptual imaging that is later "copied" in retrospective imagination. Adverbialists will likely hold that there is really just *one* basic kind of imaging process, and that it occurs more vividly in perception than in imagination. Thus, imaging blue waters is simply imaginationally, rather than perceptually, sensing in the way one does upon seeing blue waters: sensing "blue-waterly," as we might adverbially express it. Since the adverbial view conceives imaging as a way of experiencing rather than as a relation to an object, there *is* no image as an object to be copied.

On the adverbial view, then, there is no need to posit second-order images to represent first-order mental images to us, and the less vivid imagings which might seem to represent mental images are best construed as less vivid occurrences of the original imaging process. This point does not show that there *cannot* be second-order images or similar interior objects of inner vision, but the adverbial view reduces the inclination to think that there are any by suggesting a plausible alternative account of the facts needing explanation. Chief among these facts is that in recalling an image (especially a sensory image), one may have a less vivid image which apparently stands to the former as an imaginational image of a scene stands to the sensory image of that scene from which the imaginational image seems copied. The adverbial account of sensory (and other) experiences might explain this by interpreting the recalled image, say of blue waters, as *recollectively* sensing blue-waterly, where this is like visually sensing blue-waterly, but less vivid.

Given these and other points, it seems doubtful whether any realist theory of the introspection of images—one that takes them to be objects existing in their own right and having their own properties—can justify a strong analogy between that kind of introspection and ordinary viewing. For it is

by no means clear that there is any *object* introspected to serve as the counterpart of an object of ordinary vision. For an adverbial approach to experience, although realism about the (physical) objects of perception is a highly plausible view, realism about the objects of introspection is not. The idea is roughly that mental properties, such as imaging, can adequately represent physical objects in our mental life; inner objects should not be postulated for this task.

The anti-realist element in this adverbial view should not be exaggerated. To deny that mental images are objects having their own properties and in that sense are not real does not in the least imply that *imaging* is not real. Imaging processes are surely real properties of persons, even though they are not relations between persons and objects of immediate, inner perception. Does this imply, say, that introspection has no object in the sense of something it is *of* (or about), such as imaging blue waters? Certainly not. But this “object” is a kind of *content*, not an entity to which the mind is related. On the adverbial view of introspection, this kind of object is determined by what the introspection is about and is not a thing with its own properties, such as colors and shapes, sounds and movements, depths and textures.<sup>4</sup>

### *The analogy between introspection and ordinary perception*

The adverbial view in question may seem unable to do justice to the apparently causal character of introspection. There is surely some causal explanation of my being acquainted with, say, imaging blue waters rather than imaging the Statue of Liberty when I monitor a daydream of a rural summer holiday. Perhaps it is mainly in what causes the relevant imaging that such introspective consciousness differs from seeing. How might this difference be explained?

Suppose the adverbial account is true. Introspection may still be like simple perception in two ways. First, introspective viewing may imply some kind of causal relation between what is introspected in it, say an imaging, and the introspective consciousness of that state or process. Second, such viewing may imply a causal relation between the object of introspective knowledge—for instance one’s imaging blue waters—and the beliefs constituting this knowledge.

In explaining the analogy between introspection and perception, I want to concentrate mainly on introspective beliefs as compared with perceptual beliefs; we can then understand how introspection, and indeed consciousness in general, can ground justification and knowledge. A major question is how we can determine whether what the theory says is true: whether, in introspecting, as when we concentrate on our own imaging, the beliefs we thereby form about what we are concentrating on are produced by that very thing or by some aspect of it such as its imagined blue color. Only to the extent that they are should we expect introspection to ground justification and



knowledge in the broadly causal way that perception does. Many considerations are relevant here, but let me cite just two sorts.

First of all, it is surely *because* I am imagining cool blue waters that, when, with closed eyes, I introspectively consider what I am conscious of, I believe I am imagining them (and am conscious of my imagining them). It is natural (and reasonable) to take this ‘because’ to express a causal relation, even if I *could* be mistaken in thinking there is one. If the causal basis of my belief is not some inner object seen (as on the sense-datum theory), it is presumably the state or process of imagining. This is, in any event, how an adverbial theory of sensory experience would view the causal relations here. Similarly, if I introspectively believe that I am thinking about introspection, I believe this because I *am* thinking about it: it is the thinking itself that causes me to believe that it is occurring. In both cases, the introspective beliefs are produced by inner processes, and indeed in a way that makes it plausible to regard the beliefs as true. Some inner processes are *like* seeing an object in still other ways, but these processes can all be understood without positing inner objects analogous to perceptible ones such as trees and seen by the introspective eye.

A second point is this. Suppose my believing that I am imagining cool blue waters is not caused by my imagining them. The belief is then *not* introspective at all. It is *about* what is introspectable, but it is not *grounded* in introspection, any more than a belief merely about a perceptible, such as the rich red in a painting in a faraway museum, is a perceptual belief. Here, then, is another important similarity between introspection and perception.

### *Introspective beliefs, beliefs about introspectables, and fallibility*

It may seem that the case described—believing one is imagining something when in fact one is not—is impossible. But suppose I have been asked to image cool blue waters, yet I hate the water and anyway have a lot on my mind. Still, if I want to be cooperative, then even though my mind is mainly on my problems, I may call up an image. However, as I am not concentrating on calling up the image, the image that I actually get might be only of a blue surface, not of blue waters. I might now inattentively assume (and thereby come to believe) that I have called up the requested image of cool blue waters. This belief is produced by a combination of my calling up the wrong image, which I do not attentively introspect at all, and by non-imaginational factors such as my desire to cooperate. I might even retain the belief for at least some moments after I cease to image at all. In that case, it is not only not true; it is not even introspective.

This example suggests that even a true belief about one’s conscious states or processes would not be introspective without being causally connected with them. It would be about these introspectable elements but not grounded in “seeing” them in the way required for being an introspective belief. Other examples support the same point. Imagine that my task is to think about

introspection for a solid hour. I monitor myself and, on the basis of introspection, conclude from time to time that I am thinking about introspection. As I reflect on my topic, I continue to believe that I am thinking about introspection. Now when I truly believe this simply because I have repeatedly confirmed it and am confident of steady concentration, and *not* because I am still monitoring myself introspectively, my belief, though perfectly true, is not introspective.

The best explanation of this point seems, again, to be that my belief is not caused (in the right way, at least) by the thinking that should be its ground. It is a retained belief about my ongoing mental activity; it is not *grounded* in that activity as a focus of my introspective attention. My belief that I am thinking about introspection is a propositional belief *that* I am now doing so, but it is not an objectual belief, regarding my present thinking, to the effect that it is about introspection. It is not grounded in my *present* thinking at all, any more than my belief that a painting I remember portrays a rich red coat is grounded in seeing that red.

An overall conclusion we may draw here is that although there may be no objects such as sense-data or imaginational copies of them which we introspect, the process by which self-consciousness leads to introspective beliefs, and thereby to knowledge and justified beliefs about one's own mind, is nevertheless causal. Like perception of the outside world and (though in a different way) recalling events of the past, it produces something like a sensory impression and, at least commonly, beliefs about what seems to be revealed to one by that impression. The causes of introspective beliefs, however, are apparently processes and events in the mind. They are not, or at least need not be, objects that reside therein.

## Consciousness and privileged access

Suppose that introspective consciousness is causally grounded in the way we have seen. We should then raise some of the same epistemological questions about it that we raised about perception. For instance, is introspection subject to counterparts of illusion and hallucination? And if it is, how might it still be a source of justification and knowledge? Let us start with the question of how anything like illusion or hallucination might occur in consciousness.

### *Infallibility, omniscience, and privileged access*

One might think that, regarding the inner domain, which is the subject of introspective beliefs, one cannot make mistakes. If so, one might conclude that illusion and hallucination regarding this domain are impossible. Indeed, David Hume maintained that since "the contents of the mind" are known by "consciousness" (by which he meant something at least much like introspection), "they must appear in every respect what they are, and be what they appear."<sup>5</sup>

Hume's statement suggests two far-reaching claims about self-knowledge. One claim—that the contents of our minds must be what they appear to us to be—expresses the idea that introspective consciousness can give us beliefs that cannot be mistaken. The other claim—that these contents must appear to be what they are—expresses the idea that consciousness makes us so richly aware of the (introspectable) contents of the mind that it guarantees us full knowledge of them. These ideas need refinement before we can reasonably appraise them.

The first claim suggests a *thesis of infallibility* (impossibility of error): an introspectional belief—roughly one to the effect that one is (now) in an occurrent mental state (such as imaging) or that one is undergoing a mental process (such as thinking) or that one is experiencing something (such as pain)—cannot be mistaken. The infallibility thesis rests largely on the idea that we are in such a strong position regarding occurrent mental phenomena that we cannot err in thinking they are going on inside us.

The second claim suggests a *thesis of omniscience* (all-knowingness) with respect to the present occurrent contents of consciousness: if one is in an occurrent mental state, undergoing a mental process, or experiencing something, one cannot fail to know that one is. The omniscience thesis rests largely on the idea that occurrent mental phenomena are so prominent in consciousness that one cannot help knowing that they are present.

Together, these two theses constitute the *strong doctrine of privileged access*. The infallibility thesis says that our access to what is (mentally) occurring in us is so good that our beliefs about its present make-up are infallible; there is no risk of *error*. The omniscience thesis says that our access to it is so good that we cannot fail to know what (mentally) occurs in us; there is no risk of *ignorance*. It is because no one else is in such a good position to know about our mental life, and because we ourselves are not in such a good position to know about the *external* world, that it is natural to speak here of *privileged* access. The strong doctrine of privileged access is associated not only with Hume but even more with René Descartes, who is widely taken to maintain it in his famous *Meditations on First Philosophy* (1641), especially in Meditation Two.

Suppose for the sake of argument that the infallibility and omniscience theses are true. One might think that this would preclude inward counterparts of illusion and hallucination. It would not. Having illusions and hallucinations does *not* imply having false beliefs or ignorance of anything. Looking from a sharp angle from corner to corner, you can see a book as having the shape of a (non-rectangular) parallelogram without believing that it has that shape; and I can hallucinate a spruce tree like one that has burned to the ground without believing it is before me. We may know the facts. For inner perception as for ordinary sense perception, phenomenal experience is one thing and belief another.

Suppose, on the other hand, that there *are* no inner objects, such as blue, watery images, to appear to us to have properties they do not possess, such as

wavy surfaces. Then illusions of the kind we have in perception, in which an object appears to have properties it actually lacks, could not occur, as there is no object to appear to us. Nor could a hallucination of, say, an image of blue waters be *of* such an object and true or false *to* it, though of course in a sense it “misrepresents” reality. Suppose, however, that there are inner objects that we see when we image. How would hallucinating an image of, for instance, a loved one differ from just *having* that image? A sense-datum theorist might say that the hallucinatory image would be less vivid or stable than a real one. But it would still be an image of the same thing and could also be just like a normal image in vividness and other respects. It would be wrong to say, then, that a hallucinatory image is necessarily a less vivid or less stable version of a normal image, and the difficulty of explaining the difference between hallucinatory and real images is an additional reason to avoid (as the adverbial view does) positing mental images as objects.<sup>6</sup>

### *Difficulties for the thesis of privileged access*

It might be, however, that quite apart from illusion or hallucination, we can have false beliefs, or suffer some degree of ignorance, about our mental life. This is clear for *some* mental phenomena, such as dispositions like believing, wanting, and fearing. We can mistakenly believe that we do not have a certain ignoble desire (say to make a fool of a pretentious boss), particularly if it is important to our self-image that we see ourselves as having no hostile desires. For the same reasons, we can fail to know that we *do* have the desire. One can also discover a fear which, previously, one quite honestly disavowed because it was at odds with one’s sense of oneself as courageous.<sup>7</sup>

Dispositions, however, should not be conceived as *occurring*, and in any case it is the occurrent mental phenomena to which philosophers have tended to think we have the kind of privileged access expressed in the theses of infallibility and omniscience. Can we be mistaken, or at least ignorant, about our occurrent mental states or processes?

Consider first the possibility of mistake. Could I believe I am thinking about the nature of introspection when I am only daydreaming about the images and feelings I might introspect? It seems so, provided I do not attend closely to what is occurring within me. Granted, this would not be a wholesale mistake; it would be like thinking I am watching someone observing a game, when I have become preoccupied with the game itself and have ceased to attend to its observer.

Suppose, however, that the infallibility thesis is restricted to beliefs based on *attentive* introspection, where this implies “looking” closely at the relevant aspect of one’s consciousness. Call this the *restricted infallibility view*; it says only that *attentive introspectional beliefs are true*. If I carefully consider the proposition that I am thinking about introspection, and I believe it on the basis of attentive introspection (i.e. on the basis of my carefully focusing on the relevant aspect of my consciousness), could this belief be mistaken?

It may seem that error here is impossible. But suppose I desperately want to believe that I am thinking about introspection. Could this not lead me to take my daydreaming to be such thinking and even cause me to form an attentive introspective belief that I am doing such thinking? It seems so. Similarly, I could believe, on the basis of attentive but imperfect introspection, that I am imagining an octagon and then, concentrating harder and counting sides, discover that the figure has only seven.

If, for some occurrent mental states (such as thinking), it is possible to be mistaken in believing that one is now in them, then the omniscience thesis of privileged access should be abandoned along with the infallibility view. This holds even if the omniscience thesis, too, is restricted, as it should be, to cases of carefully attending to one's consciousness. The easiest way to see why fallibility cuts against omniscience is to note how omniscience would tend to guarantee *infallibility* and so would be cast in doubt if the latter is. Let me explain.

Given the extensive self-knowledge implied by omniscience, if, instead of thinking about the nature of introspection, I am only daydreaming, then I must know that I am daydreaming. But I will presumably not be so foolish as *also* to believe that I am thinking about introspection—something plainly different from daydreaming. Since I would know as well that I am occupied with, say, a series of images that portray me as swimming in cool blue waters, it is even less likely that I will believe I am thinking about introspection. It appears, then, that if I know every truth about—am omniscient about—my consciousness, then I presumably cannot believe any falsehood about it and so am infallible about it as well.<sup>8</sup>

It is at best extremely unlikely (and perhaps impossible) that these two things—knowing every truth about one's consciousness and nonetheless believing some falsehood about it—occur together, leaving one omniscient regarding one's own consciousness, yet inconsistent and fallible about it. One would know every truth about it yet would also somehow believe falsehoods incompatible with those truths. This being at best improbable, if I am fallible I am at least very likely not omniscient. Now recall our daydreaming example. It casts doubt even on the restricted thesis of omniscience. In that example, although I am in fact daydreaming, I would presumably not know that I am. If I do know that I am daydreaming, I would believe this, and then it is very doubtful that I would *also* believe I am thinking about introspection.

These points suggest that, contrary to the thesis of omniscience, I can fail to know certain things about my consciousness even when I am attending to it; but they do not imply that the omniscience side of the privileged access view is wildly mistaken, in that I might be ignorant of *every* truth about my daydreaming. Far from it. As I (objectually) believe the daydreaming to be thinking about introspection, I presumably at least know my daydreaming to involve words or colors or shapes. I have some knowledge *of* it, but I would still not know the proposition *that* I am daydreaming and thus would not be omniscient regarding the mental processes occurring in me.

*The possibility of scientific grounds for rejecting privileged access*

Perhaps there could someday be a source of significant evidence against even the restricted doctrines of privileged access. For it could turn out that every occurrent mental phenomenon is uniquely correlated with some distinct set of brain processes. Then someone could devise a “cerebroscope” for viewing the brain and could read off the contents of consciousness from the cerebroscopic data (a possibility with disturbing implications that require ethical scrutiny). What would guarantee that our introspective beliefs must match what the machine says about our mental lives? And what would a mismatch show?

Imagine that we could discover cerebroscopically a unique neural pattern for, say, believing on the basis of attentive introspection that one is imaging cool blue waters, at the same time as we discover the pattern for imaging a field of blue-green grass. It would be natural here to suppose the subject is mistaking the grassy imaging for a watery one. Might we not regard the sophisticated equipment as more likely to be right than the subject?

There is a problem with this reasoning. How could one *establish* the unique correlations except by relying on the accuracy of people’s introspective beliefs? Might it not be necessary to start by *asking* people what they are, say, imaging, to assume that they are correct, and only *then* record the associated brain state? And if learning the correlations would depend on the accuracy of introspective reports, how could the correlations show such reports to be mistaken?

A possible reply is this. First, let us suppose that learning the correlations would depend on the accuracy of introspective reports. Still, neuroscientists would not have had to rely on the accuracy of precisely the introspective belief being shown to be mistaken, and perhaps not even on the accuracy of highly similar beliefs. In any event, once they construct their instrument, they might no longer need to consult introspection to use it. They might throw away the very ladder they have climbed up on.

Imagine, however, that they do have to rely on just the sorts of belief we are examining, together with evidence regarding these beliefs’ reliability—evidence we already have independently of the cerebroscope. Would this imply that the cerebroscope could not provide powerful evidence against introspective beliefs?

Consider an analogy. We might use a mercury thermometer to construct a gas thermometer. We might calibrate a container of gas with a piston that rises and falls as the gas is heated and cooled. The new temperature readings might correlate perfectly with mercury readings in many instances: in measuring water temperature, wood temperature, and in other cases. The gas thermometer might then be used for the same jobs as the mercury thermometer *and* might gauge temperatures that the mercury thermometer cannot measure, say because they are above its boiling point. Could we not use a gas

thermometer to correct a mercury thermometer in some cases, or perhaps to correct all mercury thermometers within limits? We could. This seems so even if we originally thought the mercury thermometer infallible in measuring temperature, perhaps because we mistakenly considered its readings partly definitive of what temperature *is*. We would rebuild the ladder we have climbed up on.

Similar points might hold for beliefs about what is now occurring in one. If the analogy does extend this far—if the gas thermometer is to the mercury thermometer rather as the cerebroscope is to sincere testimony about current mental states—then even the restricted omniscience view fares no better than the restricted infallibility view. For even when I am attentive to what is in my consciousness, a cerebroscope could indicate that I do not believe (hence do not know) that a certain thing is occurring, such as a frightening image which I believe I have put out of mind.

### **Introspective consciousness as a source of justification and knowledge**

It is important not to overextend our criticism of various claims of privileged access. After all, even the restricted infallibility and omniscience views are very strong claims of privileged access. They can be given up along with the strong theses of privileged access quite consistently with holding that our access to what is occurring in us is very privileged indeed. Let us explore this.

#### *The range of introspective knowledge and justification*

Nothing I have said undermines a qualified epistemic principle. This *self-knowledge principle* says: our attentively formed introspective beliefs about what is now occurring in us are *normally* true and constitute knowledge. The difficulty of finding grounds for thinking they even *could* be false provides some reason to consider them at least very likely to be correct. Similarly, when we are attentive to what is occurring in us, then if something (knowable) is occurring in us, such as a certain melody in the mind's ear, *normally* we know that it is occurring, or at least we are in a position to know this simply by attentively forming a belief about what is going through our mind. At least this qualified epistemic principle holds for the domain of our conscious life.

Granted, our "access" to our dispositional properties is not as good as our access to what is occurring in us. We need not be conscious of the former properties, whereas the very existence of one's imaging (or of an image if there are such objects) *consists in* its place in consciousness. Beliefs and other mental dispositions need not even enter consciousness, or ever be a subject of our thoughts or concerns. Some of them may indeed be *repressed*, so that we normally cannot easily become aware of them.<sup>9</sup>

Nevertheless—and here is a *justification principle applicable to the*



*dispositional mental domain*—our beliefs to the effect that we are now in a dispositional mental state, for instance wanting, fearing, intending, or believing, are normally justified. We might also say that such beliefs, though defeasibly justified, are always *prima facie* justified, so that they are justified overall unless some defeating factor, such as an abnormal psychological interference, occurs. Moreover, normally, when we have a want (or fear, intention, belief, or similar disposition), we are in a position to know (and justifiedly believe) this. We can, then, usually know this if we need to. We very commonly do not know it, however; for such things may not enter consciousness at all, and there is often no reason to take any notice of them or form any beliefs about them. This kind of ignorance is innocuous.

There are many issues and details I have not mentioned; but we can now generalize about introspection (roughly, self-consciousness, i.e., consciousness turned inward) in relation to belief, justification, and knowledge, and summarize our main epistemological conclusions. Plainly, many beliefs arise from introspection, and the points that have emerged suggest a *general epistemic principle concerning self-knowledge* which, though far weaker than the infallibility thesis, is far-reaching: *normally, introspective beliefs grounded in attentive self-consciousness are true and constitute knowledge*. (This principle is slightly different from the self-knowledge one stated above.)

A second epistemic principle—an *attentional epistemic principle concerning self-knowledge*—though far weaker than the omniscience thesis, is that normally, if we attentively focus introspectively on something going on in us, we know that it is going on, under at least some description. I may not know that I am humming the slow movement of Beethoven's *Pathétique* Sonata, but I do know I am humming a melodic piano piece.

The corresponding justification principles suggested by our discussion seem at least equally plausible. First—to cite an *introspective justification principle*—normally, introspective beliefs grounded in attentive introspection are justified; and, second, normally, if I attentively focus on something going on in me, I am justified in believing that it is going on in me. To be sure, some such beliefs are better justified than others, and even some that are not attentive are justified. All of them are plausibly regarded as *prima facie* justified.

There are many possible principles regarding our justification and knowledge about ourselves, and there are many possible qualifications of the one just stated. But those principles are sufficient to suggest the power of introspection as a source of justification and knowledge. The examples I used to argue that introspection is fallible do not show that the apparently false introspective beliefs were *unjustified* or that true ones are not knowledge. A false belief, particularly if it is of a kind usually justified, can still be justified; and a true belief of a kind that can sometimes be false may itself constitute knowledge.<sup>10</sup>



### *The defeasibility of introspective justification*

These points about the high degree of privileged access we apparently do have may create a danger of overestimating the strength of introspective justification. From our examples, it might be thought that attentive introspection, even if not absolutely infallible, generates a kind of justification that at least cannot be defeated. Even if I am somehow mistaken about whether I am imaging blue waters, if I believe this on the basis of introspection, it would seem that I am *in the right*, even if objectively I am *not right*.

How could I fail to be justified in believing that I am imaging cool blue waters, if my belief is grounded in attentive introspection? If the question seems rhetorical, this may be because one thinks there simply is nothing else I should have done besides attending and hence no possible defeaters of my justification by appeal to another kind of ground for belief. Let us explore this.

Granting that I could not fail to be justified *unless* I could have good reason to believe I may be mistaken, still, perhaps I could fail to be justified if I had good reason for believing I am mistaken, such as evidence from repeated cerebroscopic results indicating that I have been mistaken in many quite similar cases. It is far from obvious that I could not have sufficient evidence of this sort. It seems wisest, then, to conclude that although introspective justification tends to be very strong, it remains *prima facie* rather than absolute and can be defeated by counter-evidence.

In any case, plainly beliefs grounded in attentive introspection, such as my belief that I am now imaging blue waters, are normally justified to a very high degree. Moreover—and here we have still another justification principle—*normally, my simply being engaged in attentive introspection also yields situational (propositional) justification for beliefs about what I am attending to*, even when it does not in fact yield any such beliefs. If I somehow “notice” my imaging blue waters yet do not form the belief that I am doing so, I am nonetheless (*prima facie*) justified *in* believing, and have justification for believing, that I am, just as, if I see a bird fly past and take no notice of it, I am still justified in believing it is flying past me. The analogy to perception seems sound here, and that is one reason why introspection is considered a kind of inner observation and (unless it somehow yields no content) a kind of inner perception.

### *Consciousness as a basic source*

If we now ask whether consciousness, including especially introspective consciousness, is a *basic* source of belief, justification, and knowledge, the answer should be evident. It is. In this, as in many other respects, it is like perception. But it may well be that the degree of justification which consciousness (including introspection) generates is greater than the degree generated by perceptual experience, other things being equal.

The special strength of justification on the part of beliefs about elements in consciousness has led some philosophers to think that these beliefs are a kind of foundation for knowledge and for the justification of all other beliefs. Descartes is often thought to have so regarded introspectively grounded beliefs or knowledge. Whether knowledge and justification need a kind of foundation and whether, if they do, these beliefs are the best candidates to serve as a foundation—better than, say, perceptual and memory beliefs—are the major questions pursued in Chapter 9.

There is a further epistemologically significant difference between perception and consciousness, especially as manifested in introspection, as sources of knowledge (and justification). We can by and large introspect *at will*—roughly, just by (sufficiently) wanting to—though we may also do it quite spontaneously. Moreover, there is no limit to how many things we can come to know by introspecting, if only because we can, without limit, call up images and construct thoughts. But we cannot perceive at will; and what we can know through perception is limited by what there is outside us to perceive and by external conditions of observation, just as what we can know through remembering or recalling is limited by what has actually happened (or what propositions are true) and by the conditions of belief or image retention crucial for remembering or recalling.<sup>11</sup>

Introspective consciousness, then, is unlike perception and memory in enabling us to acquire a considerable amount of knowledge whether external circumstances cooperate or not. Whatever one can “observe” in one’s own mind is a possible subject of study, and it appears that many of the beliefs we attentively form concerning our mental lives tend to constitute genuine knowledge. Very roughly, introspective consciousness is a substantially *active faculty*; perception and memory are largely *reactive faculties*.

Granted, some content—like sensations of pain—comes into consciousness uninvited. Still, we can very freely *call to mind* both propositional and imagistic content. Some of it may come from memory, which shows that introspective consciousness may draw on that as well on resources created by imagination or intellect. By contrast, sensory content, such as perceptual images, enters our mind only when our senses are *taken*, by our own observational efforts or by contingencies of experience, to it. In the inner world, by sharp contrast with the external world, there is far more at our beck and call. This is perhaps another reason why introspectively grounded beliefs have sometimes seemed to be such good material to serve as foundations for knowledge and justification. In addition to the high degree of justification self-consciousness commonly confers on beliefs, it is an active source of both justification and knowledge.

There is a trade-off, however. Through perception, we acquire (primarily) justified beliefs and knowledge about the external world; without these, we would likely not survive. Through introspection, we acquire (primarily) justified beliefs and knowledge only about the internal world; with only this, our knowledge and justification would be sadly limited, even if we could survive.

This is not to underplay the importance of the internal world: without good access to it we would have little if any self-knowledge and, for that reason, probably at best shallow knowledge of others.

Self-knowledge is also important as a back-up when questions arise about one's justification or knowledge regarding external objects. Confronted with a strange object, one may carefully consider the stability, coherence, and variations of one's perceptual experiences of it in order to rule out hallucination. Told that one merely imagined a car's passing, one may try to recall the event and then scrutinize both the vividness of one's imagery and one's confidence that the belief comes from memory rather than merely imagination. Without the kind of self-knowledge possible here, we would have less knowledge about the external world. Both perceptual and introspective knowledge are vital, and both, as we shall soon see, can be extended, by good reasoning from the raw materials they supply, far beyond their beginnings in our experience.

## Notes

- 1 To say that beliefs are dispositional properties does not imply that the concept of believing is dispositionally *analyzable*, hence equivalent to a set of conditional propositions; and I do not think the concept is dispositionally analyzable.
- 2 To be sure, images can be possessed memorially, as is my image of the Statue of Liberty when I do not have it in mind; and 'imaging' can designate a process, as when I call up the series of images corresponding to looking at the statue from the Brooklyn Heights Promenade and glancing northward to Lower Manhattan, thence to the Brooklyn Bridge, and up the East River beyond the bridge.
- 3 Both kinds of properties are experiential, in that they represent features of experience. Both, then, might be considered phenomenal, but sometimes the term 'phenomenal property' is restricted to the sensory kind that characterizes either the five senses or "inner sense," by which we feel sensations of pain and pleasure.
- 4 Such contentual objects are often called *intentional objects*, largely on the ground that, like lofty deeds we intend to perform but do not do, they need not exist.
- 5 See David Hume, *A Treatise of Human Nature* (first published in 1739–40) (Part IV, section II), ed. by L.A. Selby-Bigge (Oxford: Oxford University Press, 1888).
- 6 One might still distinguish genuine from hallucinatory images by insisting that to be a genuine image *of* (say) a loved one is to be an image caused by the corresponding sense, say seeing that person. But this view has an odd consequence. Through hearing a detailed description I could have an accurate image of Maj that is in a sense *of* her, as it perfectly "pictures" her, even if I have never seen her; but this would be mistakenly classified

as a hallucinatory image, by the causal conception just stated. There are certainly different kinds of images and various ways in which they can mislead, but the analogy between perception and introspective consciousness does not extend in any simple way to the possibility of inner illusions and hallucinations. I cannot pursue the matter further here, but for a detailed non-technical discussion of mental imagery see Alastair Hannay, *Mental Images: A Defence* (London: George Allen & Unwin, 1971) and my critical examination of this book in 'The Ontological Status of Mental Images', *Inquiry* 21 (1978), 348–61.

- 7 Some of these cases seem to occur in *self-deception*, a phenomenon that raises profound questions for both epistemology and the philosophy of mind. For a comprehensive collection of papers on it (including one offering my own account), see Brian P. McLaughlin and Amélie O. Rorty (eds.), *Perspectives on Self-Deception* (Berkeley and Los Angeles, CA: University of California Press, 1988).
- 8 The thesis of omniscience might be restricted to *introspectable* truths, as opposed to such truths as that there are 101 berries visible on the blackberry bush I am imaging, which I could know only on the basis of memory (and arithmetic) as well as introspection. The infallibility thesis might also be plausibly restricted in a similar way. This point bears on the connection between the two theses but should not affect the argumentation in the text.
- 9 Repression need not be exactly the kind of thing Sigmund Freud described, requiring psychoanalysis or special techniques to come to consciousness. There are various kinds and degrees of repression; the point here is simply that *having* a belief (or other dispositional state) is possible even if it is repressed. One might, for example, still act in the way expected of a believer of the relevant proposition. Such action is one route to discovering repressed beliefs.
- 10 For reasons to be considered in Chapter 13, skeptics tend to deny this.
- 11 There is less disanalogy in the negative cases: we cannot always cease at will to concentrate introspectively on our mental life, as illustrated by preoccupying pains; and we cannot in general cease perceiving at will; we must, for example, do so by closing our eyes or turning off a radio. This blocks the path of observation, just as an aspirin might block the path of pain.



## 5 Reason I

Understanding, insight, and intellectual power

- **Self-evident truths of reason**

*The concept of self-evidence*

*Two types of immediacy*

- **The classical view of the truths of reason**

*Analytic propositions*

*Necessary propositions*

*The analytic, the a priori, and the synthetic*

*Three types of a priori propositions*

*The empirical*

*Analytic truth, concept acquisition, and necessity*

- **The empiricist view of the truths of reason**

*Rationalism and empiricism*

*Empiricism and the genesis and confirmation of arithmetic beliefs*

*Empiricism and logical and analytic truths*

## 5 Reason I

### Understanding, insight, and intellectual power

I see the green field and I believe that it is there before me. I look away, and I believe that I am now imaging it. I remember its shape, and I believe that it is rectangular. These are beliefs grounded in my experience: perceptual, self-conscious, and memorial. But I also believe something quite different: that *if* the spruce to my left is taller than the maple to my right, then the maple is shorter than the spruce.

On what basis does one believe this obvious truth? Do we even need to see the trees to know it? Certainly it is on the basis of perception that I believe *each* of the two comparative propositions; it is easy to see, for instance, that the spruce is taller than the maple. But I do not believe on the basis of perception that *if* the spruce is taller than the maple then the maple is shorter than the spruce. As a rational being, I quite easily grasp this truth and thereby believe it.

The kind of apparently elementary use of reason this case illustrates seems basic for both knowledge and justification. But there is continuing debate about the nature and grounds of our knowledge and justification regarding the simple, obvious truths that we seem to know just in virtue of the kind of understanding of them any rational being might be expected to have. A good way to understand the epistemological role of reason is to begin with a notion that seems central for the most basic kind of knowledge and justification reason gives us—*self-evidence*.

#### Self-evident truths of reason

Such truths as the luminous one that if the spruce is taller than the maple then the maple is shorter than the spruce have been said to be evident to reason, conceived roughly as a mental capacity of understanding. They are presumably called *self-evident* because they are thought to be evidently true taken by themselves, with no need for supporting evidence. Indeed, they are often considered *obvious* in themselves, roughly in the sense that simply upon attentively coming to understand them, one normally sees their truth and thereby knows them.

### *The concept of self-evidence*

In the light of such points, we might more specifically characterize self-evident propositions as those truths such that (1) if one (adequately) understands them, then by virtue of that understanding one is justified in (hence has justification for) believing them, and (2) if one believes them on the basis of (adequately) understanding them, then one thereby knows them.<sup>1</sup> (1) says roughly that understanding them suffices for being situationally justified in believing them; it provides a justification for belief—which one can have without actually believing the proposition in question. (2) says in effect that this understanding can ground knowledge: the understanding is sufficient to render a belief based on it knowledge. (2) implies, then, that self-evident propositions are true. This implication is appropriate, since the self-evident is standardly regarded as true (and for clarity I have put truth explicitly into the characterization above).

What I have said does not imply, however, that the kind of justification one gains from understanding the self-evident is *indefeasible*, that is, so secure that it cannot be defeated, rather than *prima facie*. If the understanding in question is eliminated or obscured, the belief may cease to be justified. But at least some cases of this kind of justification are plausibly taken to exhibit justification as strong as any we can have. It can be difficult to see how defeasibility can occur here because self-evident truths are so commonly considered *also* obvious. But not all of them are—at least to finite minds. Apart from logical training, certain self-evident logical truths are not obvious to us; and it may not be obvious to most of us, on first considering it, that first cousins share a pair of grandparents. But this satisfies both (1) and (2) and is self-evident.<sup>2</sup>

There is an important analogy to perception. Just as one can see a visible property of something, such as its rectangularity, without believing that it has that property, one can comprehendingly (understandingly) consider a self-evident proposition without coming to believe that proposition; and just as one's seeing a bird fly past gives one justification for believing it did whether or not one forms this belief, adequately understanding the proposition that if the spruce is taller than the maple, the maple is shorter than the spruce, gives one (situational) justification for believing this whether one does or not.

When it comes to concepts, there is a further analogy to perception: a hierarchy analogous to the perceptual one. There is understanding a concept, such as *being taller than*. Second, there is objectually believing it to apply to something, say to a pair of things, such as the spruce and the maple. Third, there is propositionally believing something that “applies it,” as when one conceives the trees as, say, the spruce and the maple, and believes that the spruce is the taller.<sup>3</sup>

With self-evident propositions like the straightforward proposition that if the spruce is taller than the maple then the maple is shorter than the spruce,



one need not consult one's experience of the kind of thing described, or even ponder the propositions in question, in order to grasp—roughly, to understand—those propositions. And when one does come to understand them and focuses on them in the light of that understanding, one thereby normally comes to believe and know that they are true.<sup>4</sup>

### *Two types of immediacy*

There are many truths which, as just illustrated, we readily grasp and thereby immediately believe. In the simple case of comparison of heights, our belief is *immediate* in both (1) the temporal sense of 'instantly formed' and (2) the *epistemic sense*—the sense entailing that we see their truth without inferring them from anything else. The point, then, is not the temporal one that we grasp them instantly, though we may. What is crucial is that our belief exhibits *epistemic immediacy*: the belief is not based on inference or on a further, evidential belief. If it were, it would be epistemically mediate: mediated by (and thereby at least partly grounded in) the set of premises from which we infer (or on the basis of which we believe) the proposition, as my belief that Socrates is mortal is mediated by the two propositions which are part of the basis of my believing this: that he is a human being, and that all human beings are mortal.<sup>5</sup>

The proposition that Socrates is mortal is in another way unlike the proposition that if the spruce is taller than the maple then the maple is shorter than the spruce. It is not self-evident. There are at least two ways to explain why. First, Socrates and mortality are not intrinsically connected, as are one thing's being taller than a second and the second's being shorter than the first. An omnipotent God could have kept him in existence. Second (and speaking more generally), to know that Socrates is mortal one needs more than reflection—a temporally extended use of reason—on this proposition. One apparently needs information not given by the proposition. Even thinking of him as a human being does not absolutely preclude every route to his immortality. But reflection—indeed even *intuition* as a sometimes momentary use of reason—indicates that the spruce's being taller than the maple precludes the maple's not being shorter than the spruce.

This kind of point concerning propositions such as the one about the two trees has led philosophers to consider them to be *truths of reason*—roughly, truths knowable through the use of reason as opposed to reliance on sense experience. This kind of knowability has led philosophers to regard them as also *necessarily true*—*necessary*, for short, that is, such that their falsehood is absolutely precluded: there are simply no circumstances in which they are false. If a proposition is not necessary (necessarily true) *and* its negation is also not necessary, it is called *contingent*, because whether it is true—that is, its truth *or* falsity, in another terminology—is contingent on (dependent on) circumstances. That there are more than two trees in my yard is contingent. There are more, but there need not be: the number is contingent on how many I want.

## The classical view of the truths of reason

How might we understand the justification of our beliefs of self-evident and apparently necessary propositions? And how do we know them? The best-known answers to these questions, and probably the only ones we should call the *classical* answers, derive largely from Immanuel Kant, though there are similar ideas in earlier philosophers who very likely influenced Kant. He discussed both the truth of the kinds of propositions in question and how we know them.<sup>6</sup>

What Kant said is complex and difficult to interpret precisely, and I am simply going to lay out a version of the classical account which may correspond only roughly to Kant's views. Moreover, although I am interested mainly in justification and knowledge regarding the truths of reason, I will also talk about the basis of these truths themselves when that is useful in discussing how we can know or justifiedly believe them.

### *Analytic propositions*

Take the proposition that all vixens are female. I easily grasp its truth, and I immediately believe it: I depend on no premises or evidence. There was a time when 'vixen' was not in my vocabulary. I might then have looked at the sentence 'All vixens are female' and not known what proposition it expressed, much less seen the particular truth (true proposition) it does express. But this point does not show that I do not immediately believe that truth once I *do* (comprehendingly) consider it. It shows only that encountering a sentence which expresses a truth does not enable one to consider that truth unless one *understands* the sentence.

We can see, moreover, that when we do consider the truth that all vixens are female, we do not (or at least need not) know it on the basis of beliefs about the *sentence* 'All vixens are female'. For we can consider that same truth by using some other sentence to express it (say in Spanish), and perhaps without using a sentence at all.<sup>7</sup> If, however, we think about what grounds the truth of the proposition, we may discover something which in turn helps to explain why we so readily understand and believe it.

To get a sense of the ground of this truth, consider what a vixen is. It is a female fox. Indeed, the concept of a vixen may be analyzed in terms of being female and being a fox. So, in saying that a vixen is a female fox, one could be giving an elementary analysis of the concept of a vixen. Now suppose that (like Kant) we think of an analysis of a concept as indicating what the concept *contains*, or, in a certain way, *includes*. We can now say that the concept of being female is part of the concept of a vixen, and that being female is thus an element in being a vixen.<sup>8</sup>

In the light of all this, we might call the truth that all vixens are female an *analytic proposition*. To cite one major conception Kant presented, this is a proposition such that what it predicates of its subject can be "analyzed out of" the concept of that subject. Here the subject is vixens (or any arbitrarily

given vixen), and the predicate is *being female*, which is part of, and so analyzable out of, the concept of a vixen. The same sort of thing holds for the propositions that all bachelors are unmarried, that all triangles have three angles, that all sound arguments have true premises and true conclusions, and so on. Analytic propositions are usually considered clear cases of the self-evident.<sup>9</sup> This is explainable in terms of our account of the self-evident—provided we make the not implausible assumption that, given an adequate understanding of such a proposition, one can frame an analysis in which the containment relation is clearly evident. By contrast, Kant and others have viewed non-analytic propositions as *empirical*, and taken empirical propositions to be knowable not by using reason alone but only on the basis of confirmatory experience—most prominently (and perhaps necessarily) perception.

### *Necessary propositions*

This way of looking at our example helps to explain something else that is true of the proposition that all vixens are female: it *cannot* be false and, in that sense, is necessary (a necessary truth). To see this point, try to conceive of a non-female vixen. Since the concept of a vixen is analyzable as (and hence equivalent to) that of a female fox, one is in effect trying to conceive of a *non-female* female fox. This would be both female and not female. We would have a contradiction. Hence, there cannot be such a thing, on pain of contradiction. It is thus absolutely impossible—in a sense implying impossibility by the laws of logic—that there be a non-female vixen. By contrast, it is possible that there is, and also that there is not, a 200-pound vixen. The proposition that all vixens weigh less (or more) than this is *contingent*: neither necessarily true nor necessarily false.

Because the falsity of analytic propositions entails a contradiction in this way, they are often thought to be—and are sometimes even *defined* as—those that are true *on pain of contradiction*. That is, their falsity entails a contradiction, and hence they can be false only if a contradiction is true. That is absolutely impossible. Analytic propositions are therefore regarded as truths that hold in *any* possible situation and hence are necessary (though other kinds of truths are also considered necessary).

Now if analytic propositions are true by virtue of the sort of conceptual containment relation we have been exploring, might we not know each one we do know in virtue of grasping the containment relation basic to it, in the sense that we have an adequate understanding of that relation? In considering the proposition that all vixens are female, one in some way grasps the containment relation between the concept of a vixen and that of being female. Intellectually—*intuitively*, in one widely used terminology—one sees the relation and thereby sees and (non-inferentially) knows the truth it underlies.

It might be objected that the correct account is instead this. One quickly or subconsciously reasons: The concept of a vixen is analyzable as that of a female fox; *being female* is contained in that analysis; hence all vixens are

female. So, it may be claimed, one knows that all vixens are female only inferentially. A defender of the classical view would reply that this second-order reasoning indicates how one might *show* that one knows that all vixens are female, but it does not indicate *how* one *knows* it, at least not if one just grasps its truth in the normal way.

The classical account can grant that one perhaps *could* come to know the proposition in that indirect way, by conceptual analysis. But one need not come to know it in that way; and normally, if one did not *already* know that vixens are female foxes, one would not even be in a *position* to know (on one's own) the sophisticated truth that the *concept* of a vixen is analyzable as that of a female fox. Believing that all vixens are female, in virtue of grasping the crucial containment relation between the concept of a vixen and that of a female, does not require coming to know this proposition in that sophisticated way.

### *The analytic, the a priori, and the synthetic*

We can now see how the classical account of the truths of reason might apply to apparently *non-analytic* truths that are directly and intuitively grasped. Think about the proposition that nothing is both red and green all over at one time (different kinds of examples will be considered in Chapters 12 and 14). This is apparently self-evident and hence a truth of reason. But is it analytic? Can we analyze *being non-red* out of the concept of being green, or *being non-green* out of the concept of being red, so that anyone who said that something is red and green all over at once could be shown to be implying that it is (wholly) red and non-red, or green and non-green? This is doubtful. For one thing, it is not clear that we can analyze the concept of being red (or the concept of being green) *at all* in the relevant sense of 'analyze'. Still, on the classical view, we can know through the use of reason the necessary truth that nothing is red and green all over at once.

Let us consider two kinds of objections to the idea that the proposition that nothing is red and green all over at once is self-evident and necessary, yet not analytic. The first is based on treating the proposition as empirical and contingent; the second objection says it is analytic after all.

Take the contingency objection first. One might think that there could be a scientific explanation of why nothing is red and green all over at once; and if there is, then (on a plausible and standard view of such matters) the proposition is empirical and not self-evident or even necessary. How might such an explanation go? We can, after all, scientifically clarify what being red (or any other color) is by appeal to facts about light. This might seem to enable us to know all there is to know about basic relations among colors, even though the relevant facts about light are contingent. On the classical view, however, although scientific investigation helps us to understand certain facts about red *things* (and perhaps about the property of being red), it does not indicate what is essential to the *concept* of a red thing, such as being non-green at

the time it is red. Similarly, it is essential to the concept of a vixen that it is equivalent to that of a female fox.

To be sure, one could discover scientifically that vixens have a unique tracking system. But normally one would be identifying them for study as *female* foxes and hence would not set out to discover whether they are female. On the classical view, we cannot identify anything as a vixen—say, for experimental purposes—except under the assumption that it is female. Thus, the possibility of discovering anything inconsistent with its being female is ruled out from the start. If our experimental subject is *selected* by its having a specified property, we cannot find out experimentally that *it* (as opposed to something else it may turn into) lacks that property.

Similarly, one would not normally set out to discover scientifically whether what is red all over is ever also green all over at the same time—since it would be at best difficult to wonder whether this is true without immediately *seeing* that it is. This does not make analytic or any self-evident truths more important than scientific truths. The former are simply different: they are not of the right kind to be open to scientific verification or falsification, and in part for this reason they also do not compete with scientific truths.

It appears, then, that the suggested “scientific” objection to the classical view fails. If, however, the proposition that nothing is red and green all over at once is not a “scientific truth,” that might be because it is analytic after all. Let us explore further whether the classical view is correct in claiming that the two self-evident truths in question still differ in this: being non-green is not analyzable out of the concept of being red, whereas being female is analyzable out of the concept of being a vixen.<sup>10</sup>

This brings us to the second objection. The objection proceeds by arguing (against the classical view) that the proposition that nothing is red and green all over at once is analytic. Could one not *indirectly* analyze the concept of being red as equivalent to the concept of having a color *other than* green and blue and yellow, and so on, in which we list all the remaining colors? This claim may seem right, because it *seems* self-evidently true that red is the only color filling that bill. But the claim is doubtful. For one thing, it is questionable whether a determinate list of all the other colors is even possible. More important, even if it is possible, the concept of being red is not *negative* in this way. To be red is to have *that color*; to be red is not simply to be a color other than green, blue, yellow, etc. Third, there is an important disanalogy to our paradigm of the analytic: whereas one could not have the concept of a vixen without having the concepts of a fox and a female, one could have the concept of being red (and so understand that concept) without even having all of these other color concepts (even if one must have *some* other color concept).

Moreover, proponents of the classical view would stress here (what is independently plausible) that an analysis does not merely provide a *conceptual equivalent*, that is, one which (necessarily) applies to the same things to which the concept being analyzed does, as the concept of being not-not-red applies

to everything the concept of being red does. An analysis of a concept (as we shall see in Chapter 11 in exploring analyses of the concept of knowledge) must meet at least two further conditions. First, it must exhibit a suitable subset of the elements that constitute the concept; second, it must do so in such a way that one's seeing that they constitute it can yield some significant degree of *understanding* of the concept. The concept of being red is surely not constituted by the complex and mainly negative property of being a color that is not green, not blue, and so on; and one could not understand what it *is* for something to be red simply in terms of understanding that long and perhaps indefinite list.

The relevant notion of understanding is *understanding of*, not *understanding that*, which is plausibly considered a special kind of knowledge of the proposition indicated by the 'that' clause, as in the case of understanding that citizenship requires being politically informed. Understanding *of* has an intimate connection with explanation. The implication of this point here is that an analysis of a concept must provide sufficient understanding of it to provide at least some *explanation* of it. The analysis of the concept of a vixen as a female fox provides material for an elementary explanation of that concept; but noting that being red is equivalent to being non-green, non-blue, and so on for all the other colors would not provide any explanation of what it is to be red. The concept of being red is simply not thus analyzable. Even the property of being red is not identical with that negative property. Indeed, one could presumably understand the list of other colors quite well even if one had never seen or imagined redness, and one had *no* perceptual, imaginal, or other concept of redness.

The point that an analysis should provide understanding of the kind that goes with explanation must not be taken to imply that we can have understanding *only* when we can explain. It is arguable, in fact, that the concept of redness is *simple* in the sense that, unlike that of a vixen, it is not analyzable into elements of any kind. One's understanding of the concept does not require its analyzability; it is enough to be able (above all) to apply it to the right things, withhold it from the wrong ones, and see what follows from its application—such as the thing's not being green.

On balance, then, it appears that the proposition that nothing is red and green all over at once is not analytic. This does not, however, prevent our rationally grasping the truth of that proposition. Truths that meet this rational graspability condition—roughly a knowability through conceptual understanding condition—have been called *a priori propositions* (propositions knowable 'from the first'), because they have been thought to be such that they can be known *a priori*, in a very strict sense of this phrase: known not on the basis of sense experience but simply through reason as directed toward them and toward the concepts occurring in them, at least if reason is used extensively enough and with sufficient care. Propositions that are *a priori* in this strict, knowability sense—as is the proposition that nothing is red and green all over at once—are also plausibly considered self-evident.<sup>11</sup>

Moreover, the kind of justification for believing a self-evident proposition when we believe it in the indicated way is a basic kind of justification and is often called a priori.

By contrast with analytic propositions, however, the kind of a priori proposition exemplified by that one seems to assert something beyond what analysis of the relevant concepts can show. For this reason, propositions of this kind are also called *synthetic propositions*, though these are typically defined negatively, simply as *non-analytic*. Positively conceived, they typically bring together or “synthesize” concepts and properties, even if in a negative way (as by linking redness with colors other than green—by including it among these other colors). Synthetic propositions need not, even *in part*, analyze concepts, and many are empirical in the straightforward way in which propositions evident to the five senses are.

It is noteworthy that although analytic propositions are characterized roughly in terms of how they are *true*—by virtue of conceptual containment (or, on a related account, on pain of contradiction)—a priori propositions are characterized in terms of *how they are known*, or can be known: through the operation of reason.<sup>12</sup> (This allows that they can *also* be known through experience, say through receiving testimony, at least if the attester’s knowledge is, directly or indirectly, grounded in the operation of reason.)

On this basis, a priori propositions are also negatively characterized as knowable “independently of experience,” in which this phrase above all designates no need for evidential dependence on experiential grounds, such as those of perception. But even if this negative characterization of a priori propositions is correct so far as it goes, understanding them through it will require understanding of the kinds of positive characteristics I am stressing. Let us pursue these.

### *Three types of a priori propositions*

If we take knowability through the use of reason as a rough indication of what constitutes the a priori in general, then it includes not only self-evident propositions but certain others that are not self-evident: most clearly those propositions *not* themselves knowable simply through reason as directed toward them and toward the concepts occurring in them, but self-evidently *following from* (entailed by) such (self-evident) propositions. This is the simplest case of what is *a priori in the broad sense*. Consider the proposition that either nothing is red and green all over at once or I am flying to the moon. This self-evidently follows from the proposition about red and green, which (apparently) is self-evident. It self-evidently *follows* because it is self-evident that *if* nothing is red and green all over at once, then *either* that is true *or* I am flying to the moon.

One might think that this disjunctive (either–or) proposition is self-evident because it is so obviously both true and necessary. But even though this is true, one knows it, not in virtue of understanding the proposition



itself, but in virtue of its self-evidently following from something that *is* self-evident. One knows it inferentially, on the basis of knowing the simpler proposition that nothing is red and green all over at once. One cannot know it *just* from understanding it, as with a self-evident proposition, but only through seeing the quite different truth that if nothing is both red and green at once, then either that proposition is true or I am flying to the moon. This conditional (if–then) proposition *is* self-evident; hence, it is an utterly secure ladder on which to climb from knowledge that nothing is red and green all over at once to knowledge that either this is so or I am flying to the moon. That disjunctive proposition is *a priori in the broad sense*.

Suppose, however, that a proposition is neither self-evident nor self-evidently entailed by a self-evident proposition, but *is* provable by self-evident steps (perhaps many) from a self-evident proposition. Because there is more than one step and there can be many steps, such a provable proposition might or might not be knowable without reliance on memory, depending on the mental capacity of the rational being in question. Nonetheless, since it can be known through such a *rigorous proof*—one that begins with a self-evident proposition and proceeds only by self-evident steps (entailments) to its conclusion—a rigorously provable proposition may be called *ultimately a priori* (or ultimately self-evident, though the former term seems preferable). It is not *a priori* in the broad sense because (1) it is not linked to the self-evident by a *single* step and—more important—(2) it is not necessarily self-evidently linked to it.<sup>13</sup> But as it is ultimately traceable to a self-evident proposition, it may be considered *a priori* in the ultimate provability sense.

Thus, in speaking of propositions that are *a priori* in the most comprehensive terminology, I include not only the intuitively central cases that are self-evident or just one step from it—propositions self-evidently entailed by a self-evident proposition—but also those not thus entailed but nonetheless provable by self-evident steps from a self-evident proposition.

We could say, then, that for the kind of classical view in question, the self-evident is the *base* of the *a priori*: *a priori* propositions are those that are either self-evident (i.e., *a priori* in the narrow sense) or, though not themselves self-evident, self-evidently follow from at least one proposition that is (hence are *a priori* in the broad sense). The general notion of an *a priori* proposition, applicable to both cases and others, is roughly the notion of a truth that either is a self-evident proposition or is self-evidently entailed by one, or provable from one by self-evident steps.<sup>14</sup>

Knowledge of propositions *a priori* in the broad or ultimate provability sense, unlike knowledge of those *a priori* in the narrow sense, depends on knowledge of some self-evident proposition as a ground. But neither kind of knowledge depends on knowledge of any empirical proposition, and in that sense both kinds are “independent of experience.”

It is because *a priori* propositions (of any sort) are understood in relation to how they can be *known* that the notion of the *a priori* is commonly considered epistemological. But many *a priori* propositions also have a special



property of a different kind. Many are said to be analytic. The notion of the analytic is more often taken to be of a different, non-epistemological kind, say conceptual, since analytic truths are conceived as grounded in a simple containment relation between concepts.<sup>15</sup>

It should not be surprising, then, that the categories of the analytic and the *a priori* are not identical. In both cases, however, proponents of the classical view have taken the relevant propositions to be necessary: this is commonly thought to be obvious for the analytic ones, which are true “on pain of contradiction,” but it has seemed reasonable to classical theorists to hold that even synthetic *a priori* propositions are invariably necessary. Perhaps the underlying thought is that if their truth were contingent and so depended on what holds in (is contingent on) some possible situations but not others, one could not know it just on the basis of understanding the proposition itself.

### *The empirical*

A huge variety of truths are not *a priori*. That the spruce is taller than the maple is one of them. Truths that are not *a priori* are called *empirical* (or *a posteriori*) truths. This means, roughly, that the propositions in question can be known only *empirically*: knowable (assuming they are knowable) only on the basis of experience, as opposed to reason—above all on the basis of perceptual or self-conscious experience (in the ways described in Chapters 1, 2, and 4).

Saying simply that a proposition is empirical (or *a posteriori*) leaves open whether it is true: there are empirical falsehoods, such as that it is not the case that the spruce is taller than the maple, as well as empirical truths. (In this the term ‘empirical proposition’ is unlike ‘*a priori* proposition’ and ‘necessary proposition’, which are not commonly used to refer to falsehoods, but my main examples of empirical propositions will be truths.)

For the classical view, empirical propositions as well as *a priori* propositions are crucial for our lives. Indeed, the former include every truth known perceptually, such as those known through observing the colors and shapes of things, and all truths known scientifically, such as generalizations linking the temperatures and the volumes of gases, or ingestion of drugs with change in behavior. A certain range of *a priori* propositions, such as those of logic and pure mathematics, are presupposed by both common sense and science. Empirical propositions are also required to guide us in dealing with the world, but the classical view sees them as open to disconfirmation through experience in a way that *a priori* propositions are not.

### *Analytic truth, concept acquisition, and necessity*

Analytic truths, as well as certain synthetic ones, are called *a priori* because analytic truths are knowable through the use of reason. But analytic truths appear to be knowable—or at least are showable—through a different use

of reason than is appropriate to the synthetic a priori truths. It may be that I know that nothing is red and green all over at once by virtue of simply grasping, as a rational being, a kind of incompatibility between the concept of being red (at a time and place) and the concept of being green. But, as pointed out earlier, I apparently do not know it by virtue of grasping a *containment* relation between being red (or green) and anything else. If this does not illustrate two different uses of reason, it at least indicates a different kind of application of reason to different kinds of relations of concepts.

Because my knowledge of the proposition that nothing is red and green all over at once is not based on grasping a containment relation, it differs from my knowledge of the analytic truth that all vixens are female. Yet in both cases the relation between the concepts involved in the truth seems to be the basis of that truth. In both, moreover, I apparently know the truth through rationally *understanding* that relation: a relation of analytic containment in one case, and of mutual exclusion in the other.

These points do not imply that experience is irrelevant to knowledge of the a priori. On the classical view, I do need experience to *acquire* the concepts in question, for instance to acquire color concepts or the concept of a fox. But once I have the needed concepts, it is my grasp of their relations, and not whatever experience I needed to acquire the concepts, which is the *basis* of my knowledge of analytic and other a priori truths.

In part because of these similarities, as well as because the falsity of a priori propositions seems inconceivable, the classical view takes synthetic a priori truths as well as analytic truths to be necessary. They cannot be false, even though in the synthetic a priori cases it seems not to be strictly contradictory to deny one. For instance, claiming that something *is* red and green all over is not contradictory in the sense that it (formally) entails that some proposition—say, that the object in question has a definite color—is and is not true. Still, on the classical view it is absolutely impossible that something *be* red and green all over at once. We need only reflect on the relevant concepts (mainly the color concepts) to realize that nothing is red and green all over at once; we readily grasp (apprehend) an exclusion relation between being red and being green.

It is also commonly held by philosophers in the classical tradition that all necessary propositions are a priori. One rationale for this might be that necessity is grounded in relations of concepts and these (or at least the relevant relations) are the same in all possible situations. A mind that could adequately survey all possible situations (like the divine mind as often conceived) could thus know the truth of all necessarily true propositions. Since this survey method would be possible without analyzing one concept out of another, the grounding of necessity in conceptual relations would also explain how there can be synthetic necessary truths. And for the classical view, these, being necessary, are also a priori.<sup>16</sup>

Summarizing, then, the classical view says that all necessary propositions are a priori and vice versa, but it maintains that analytic propositions

constitute a subclass of a priori ones, since some a priori propositions are synthetic rather than analytic. The view tends to conceive the truth of all a priori propositions as grounded in relations of concepts (or of similar abstract entities, such as “universals,” in Bertrand Russell’s terminology).<sup>17</sup> But the position conceptually accounts for these propositions differently: for necessary propositions in terms of the unrestricted *circumstances* of their truth (the absolute impossibility of their falsehood in any circumstances), for analytic ones in terms of *how* they are true (typically, by virtue of containment relations), and for a priori propositions in terms of how their truth is *known* (through understanding).

### The empiricist view of the truths of reason

The classical view of the nature of what I am calling *a priori truths*—*truths of reason*—and of our knowledge of them has been vigorously challenged. To appreciate the epistemological significance of reason as a source of justification and knowledge, and of truths of reason themselves, we must consider some alternative accounts of these truths.

John Stuart Mill held that ultimately there are only empirical truths and that our knowledge of them is based on experience, for instance on perception.<sup>18</sup> We might call this sort of view *empiricism about the (apparent) truths of reason*. The name suits the view, since the position construes apparently a priori truths as empirical, though it need not deny that reason as a capacity distinct from perception has *some* role in giving us justification and knowledge. Reason may, for example, be crucial in extending knowledge by enabling us to prove geometrical theorems from axioms. But the view I want to explore (without following Mill in particular) denies that reason grounds justification or knowledge in the non-empirical way described by the classical theory.

### Rationalism and empiricism

Before we consider Mill’s thesis in detail, we should contrast it, from the most general epistemological point of view, with that of Kant and other rationalists to get a better sense of what is at stake in the controversy between rationalism and empiricism. Kant’s position on the truths of reason might be called rationalist, Mill’s empiricist. These terms are used too variously to make precise definition wise. Very roughly, however, *rationalism* in epistemology takes reason to be far more important in grounding our knowledge than empiricism allows, and rationalists virtually always assert or imply that, in addition to knowledge of analytic truths, there is knowledge of synthetic a priori truths. Very roughly, *empiricism* in epistemology takes experience, most notably sensory experience, to be the basis of all of our knowledge except possibly that of analytic propositions, understood as including purely logical truths, such as the truth that if all whales are mammals and no fish

are mammals then no whales are fish. (For both empiricists and rationalists, analytic propositions are typically taken to include logical truths.)<sup>19</sup>

One might wonder why some empiricists grant that analytic truths may be a priori. The central point (though an empiricist might not put it this way) may be seen if we use the terminology of the classical theory: even if such logical propositions are not true by virtue of containment relations between concepts, their negations formally entail contradictions, for instance that some vixens are and are not female foxes. They are therefore paradigms of truths of reason; for the use of logic alone, which is perhaps the purest use of reason, can show that they can be false only if a contradiction is true—which is absolutely impossible. This is another reason why, as noted above, analytic propositions are sometimes given a broader characterization than I have proposed and are taken to be those whose negations entail a contradiction.<sup>20</sup>

Some empiricists do not allow that any knowledge, even of so-called analytic propositions, is genuinely a priori. A *radical empiricist*, such as Mill, takes *all* knowledge to be grounded in experience. A *radical rationalist* (which Kant was not) would take all knowledge to be grounded in reason, for instance to be intuitively grounded in a grasp of self-evident propositions or deductively based on inference from a priori truths that are intuited.<sup>21</sup>

### *Empiricism and the genesis and confirmation of arithmetic beliefs*

Empiricism about what are called the truths of reason is most plausible for the apparently synthetic a priori ones, so let us sketch it with reference to an apparently synthetic kind of a priori proposition that has been much in dispute. Mathematical truths, particularly truths of simple arithmetic, are often regarded as synthetic a priori. Consider the proposition that  $7 + 5 = 12$  (Kant's example, also found in Plato's *Theaetetus*). It is easy to say that one just knows this, as one knows that nothing is red and green all over at once. But how does one know it?

Here we cannot readily find a good analogy for the simple exclusion relation we apparently grasp in the case of red and green. Could it be that from experience with objects, say with counting apples, then combining two sets of them, and recounting, we learn our first arithmetic truths and then use reason to formulate general rules, such as those for calculating larger sums?

Viewed in this way, arithmetic develops rather as scientific hypotheses are often thought to, with observations crucial at the base, generalizations formulated to account for them, and broader generalizations postulated to link all the observations and the narrower generalizations together. And do we not first learn to add by counting physical things, or by counting on our fingers?<sup>22</sup>

To be sure, we perhaps cannot imagine how the number 7 added to the number 5 could fail to equal the number 12. But that is not a point about the behavior of objects in the physical world. The physical world *could* go

haywire so that when (for instance) five apples and seven oranges are physically combined, the result of counting the new set is *always* eleven. If that happened pervasively, might we not begin to think that arithmetic must be revised, just as Einstein's work showed that the physics of "the incomparable Sir Isaac Newton" needed revision? Perhaps the crucial epistemological consideration is what overall account of our experience is most reasonable; and if the best overall account should require rejecting a proposition now considered a priori and necessary, so be it.

The classical view provides for several critical responses. One concerns the distinction between two quite different things: the *genesis* of one's beliefs—what produces them—and their *justification*, in the sense of what justifies them. A second point concerns whether arithmetical propositions can be tested observationally. A third focuses on the possibility of taking account of what looks like evidence against arithmetical truths, so that even if one's final epistemological standard for judging a proposition is its serving the demands of the best overall account of experience, these truths can be preserved in *any* adequate account. Consider these ideas in turn.

First, granting for the sake of argument that our arithmetic beliefs arise from counting physical objects, is the experience that produces them what *justifies* them? The genesis of a belief—what produces it—is often different from what justifies it. The testimony of someone I realize is unreliable might, when I am off guard, produce my belief that different brands of aspirin do not, apart from additives, differ chemically. My belief would at that point be unjustified; but it might become justified later when I learn that aspirin is simply acetylsalicylic acid. Moreover, regardless of what produces our arithmetic beliefs initially, when they are justified in the *way* my belief that  $7 + 5 = 12$  now is, experience does not appear to be what justifies them. For my part, I do not see precisely how the truth of the proposition might be grounded in the behavior of objects when they are combined; and I would not try to justify it, as opposed to *illustrating* it, by citing such behavior.

This brings us to the second point: it is doubtful that the proposition that  $7 + 5 = 12$  is (empirically) *testable*, say by examining how objects combine, though it is *exemplifiable* in that way. The empiricist might reply that this by no means shows that the proposition is, as the classical view insists, necessarily true rather than contingent and empirical. Indeed, it does not. But let us look closely at the idea that it could be tested, and could thereby be disconfirmed by discovering that when groups of five objects are combined with groups of seven, we find just eleven.

This brings us to a third response. How might one deal with repeated and systematic counter-evidence? Classical theorists will argue that it is possible for the world to alter in such a way that this combination procedure results in one item's disappearing, or in our failing to see it, or in our misremembering how many items entered the mix before our re-counting. They will also argue that the unexpected realization of such possibilities would be a better interpretation of the strange cases described—hence of our overall

experience—than saying that it has turned out to be false that  $7 + 5 = 12$ . Thus, instead of saying that an arithmetical principle has been falsified, we would say that the world no longer uniformly exemplifies it.

One consideration favoring the classical view is that it is at best difficult even to understand how the purely arithmetical principle could be false. The number 7 plus the number 5 apparently equals the number 12, regardless of how apples and oranges behave. For the arithmetic statement is apparently not *about* apples and oranges, though (so far as we know) their behavior exemplifies it. For the classical view, at least, it is about *numbers*, which, unlike the arabic or roman or other *numerals* we use to represent them linguistically, are abstract and non-physical. If a proposition is not about concrete objects, facts about their behavior are not a test of its truth.

Notice something else. In order to gather purportedly significant counter-evidence to the arithmetic proposition in question, one would have to rely, as already noted, not only on memory and perception (both highly fallible sources) but also on simple arithmetic: one would have to *count* disconfirming cases. A single apparent instance, say of seven and five things brought together and not adding up to twelve, would not be significant, and one must keep track of how many anomalies there are, relative to confirmatory instances in which the expected sum is counted out. It is not normally reasonable to give up a good theory on discovering a single apparent counter-instance. It appears, then, that in order to take seriously empirical evidence that would undermine arithmetic, we must trust perception in our counting, arithmetic itself in our summing, and memory in our overall judgment.

One might think it is enough simply to *have* a significant number of such disconfirming cases. But this is not so. One must be justified in believing that the number *is* significant. And how could one achieve this if one either made no count or—in any case—could not rely on one's count of single cases to sum to a significantly large number? If it need not be true that  $7 + 5 = 12$ , why should  $1 + 1 + 1$  disconfirming instances necessarily sum to 3? And would anything less than a huge number of apparently disconfirming cases be evidentially decisive against such a proposition of simple arithmetic? A single disconfirming instance would surely seem just an anomaly; there must be a significant number. One would, then, have to rely on some arithmetic propositions, such as that  $1 + 1 + 1$  disconfirmations = 3 (a minimally significant number, perhaps), in order to mount an effective challenge to the (necessary) truth that  $7 + 5 = 12$ . Given the interconnections among arithmetic propositions, it is not clear that one could consistently (or at least with any plausibility) maintain the needed disconfirmatory propositions while denying that  $7 + 5 = 12$ . Still another obstacle to recognizing apparent counter-evidence as genuine is the dependence on memory to keep track of disconfirming instances. The fallibility of memory would defeat confidence that one had adequately tracked apparent disconfirmations.

There may be a way around these objections, but even finding it would leave one far from a strong case for the contingent or empirical status of

arithmetic truths.<sup>23</sup> Even if one appealed, not to apparent counter-instances to the proposition that  $7 + 5 = 12$ , but to a well-confirmed *theory* to argue that it could be false, one would need to do at least some counting of one's confirmatory data regarding that theory (not to mention other ways in which theory confirmation relies on arithmetic, perception, and memory).

None of these points requires us to deny that there *is* a similar, contingent arithmetic proposition about apples and oranges, namely that when we *count* five of the first and *place* them next to the result of counting seven of the second, we *can count* twelve all told. This proposition may easily be confused with its pure mathematical counterpart. The former is clearly contingent and empirical, but its being so does not show that the purely arithmetic proposition is also. The distinction between pure and applied mathematics can also be brought to bear on geometry.<sup>24</sup>

There is a related metaphysical dimension of the question of the status of arithmetic truths. By contrast with at least one form of the classical view, radical empiricism denies that there *are* abstract entities and so, believing that mathematical propositions are about something concrete, radical empiricists naturally view them as generalizations about the behavior of physical objects. We need not accept the empiricist view to grant that if physical things did not exemplify the proposition that  $7 + 5 = 12$ , the proposition would be of far less *value* to us even if necessarily true. If the physical world went haywire, it could turn out to be false that when seven apples are placed together with five more and the total collection is counted, the count yields twelve. This chaotic situation would falsify the *physical* principle already contrasted with the arithmetic one in question. But the physical principle is not, and does not even follow from, the purely mathematical proposition we are discussing.

### *Empiricism and logical and analytic truths*

The empiricist view of the a priori can also be applied to analytic propositions and even to self-evident logical truths, and it may indeed appear more plausible in that case. Suppose that through scientific investigation we discover that vixens have certain characteristics we think of as male, such as certain hormones. Imagine that gradually (perhaps because of chemicals in the environment) these discoveries mount up so that the female foxes in our laboratory begin to seem more aptly classified as male than as female. Could not a time come when we begin to doubt that vixens are female after all?

And what about the logical principle of the excluded middle, which says that every proposition is either true or false? Consider the proposition that Tom is bald. Must this proposition be either true or false no matter what the quantity or distribution of hair on his head? Surely the proposition is an appropriate counter-example to the principle of the excluded middle.<sup>25</sup>

The classical view can offer its own account of these examples. For one thing, particularly over a long time, we can begin to use a term in a sense different from the one it now has. Thus, the discoveries about vixens could result in our someday using 'vixen' to mean not 'female fox', but 'fox with



female external sexual characteristics and of the anatomical kind *K*' (where *K* is the kind of animal we have in our laboratory). Then, when we utter such words as 'Vixens are not really female', we are not denying the analytic proposition *now* expressed by 'All vixens are female'. We have confirmed something else, rather than disconfirming this.

In this way, then, our experience might result in our someday no longer assertively uttering 'Vixens are female' to say anything that we believe. This certainly does not show that experience might falsify the proposition we *now* affirm when we assertively utter that. Given what we now mean by 'vixen', in saying that all vixens are female we do not rule out that *those* 'vixens' in the lab could have internal biological and chemical characteristics in the light of which *they* ultimately need not be considered female.

Regarding the principle of the excluded middle, I would stress that Aristotle plausibly argued against it, and some contemporary philosophers of logic do, too. The main reasons for doubting it, moreover, do not depend on empiricism. Let us explore some of them.

Consider again the vague statement that Tom is bald. It may certainly be argued that this need not be either true or false. It is not as if 'bald' meant, say, 'having fewer than 500 hairs on the top of one's head'. It does not. And if it did, the term 'top' would still be vague and would cause the same trouble: it would be unclear in what *area* we must find 500 hairs. If the middle possibility—neither truth nor falsity—is to be ruled out, it must be by a better argument. The principle of the excluded middle, though often used to suggest that even logical truths are not necessarily true, is controversial among rationalists and empiricists alike. The principle is a poor example to support the empiricist case against the necessity of logical truths.

When, by contrast, standard examples of simple logical truths are used, the effect seems very different. Consider the proposition that if Ann is coming by bus or she is coming by plane, and it is false that she is coming by bus, then she is coming by plane (which exemplifies the general logical truth that if at least one of two propositions is the case and the first is false, then the second is true). Is there any plausibility in the view that this might be false? I find none; and while nothing said here proves that the empiricist account of the *a priori* is mistaken, it appears less plausible than the classical account.

If what we have seen so far is accepted, the classical view of the truths of reason is quite defensible and the empiricist critique of it fails. But we have not yet adequately taken into account the ways in which knowledge of those truths might depend on language. This is an important topic particularly given the extent to which understanding the *a priori* is connected with understanding language. The next chapter will consider this topic in some detail.

## Notes

- 1 *Adequacy* of understanding of a proposition cannot be merely partial understanding, and it is more than simply getting the general sense of a sentence expressing it, as where one can analyze the grammar of the



sentence, indicate something of what it means through examples, and perhaps translate it into another language one knows well. Adequacy here implies not only seeing what the proposition says but also being able to apply it to (and withhold its application from) an appropriately wide range of cases. This matter is treated in some detail in my 'Self-Evidence', *Philosophical Perspectives* 13 (1999), 205–28. Note also that there is no appeal here to understanding or positing the *necessity* of the propositions (though the characterization lends itself to taking them to be necessary). In this respect my notion of the self-evident is simpler and more moderate than the traditional one common in much of the literature. See, for example, Laurence Bonjour, 'Toward a Moderate Rationalism', *Philosophical Topics* 23, 1 (1995), 47–78, esp. section 3.

- 2 For a helpful discussion of obviousness related to (but quite different from) the one in my 'Self-Evidence' and connected with the theory of the a priori in general, see Robin Jeshion, 'On the Obvious', *Philosophy and Phenomenological Research* 60, 2 (2000), 333–55.
- 3 Two points are appropriate here. (1) A fourth case is one in which a concept is not only *exercised* in a belief but explicitly figures in it, as when one believes that the concept *being taller than* is instantiated by the spruce and the maple. (2) The analogy between perception and conception I am developing is meant to leave open what concepts are and what it is to understand one. As will later be apparent, philosophers differ in their understanding of the truths of reason in part because of their different understandings of the nature of concepts.
- 4 One reason for the normality qualification is to make room for the possibility that one can consider and adequately understand a self-evident proposition yet fail to believe it. Brain manipulation might cause such failure. We should also make room for the possibility that, especially with more complex self-evident propositions—say that if  $p$  entails  $q$  and  $q$  entails  $r$  and  $r$  entails  $s$ , and  $s$  is not true, then  $p$  is false—it may take a person time to form the belief.
- 5 *Temporal immediacy*, unlike epistemic immediacy, is a property not primarily of beliefs as such but of their formation. A belief is temporally immediate when its formation occurs "without delay" upon the person's considering the proposition in question (or encountering the situation, such as the sight of a lightening bolt, that gives rise to the belief). One could also say that propositions are temporally immediate in a derivative sense when they are so obvious that one normally believes them immediately on (comprehendingly) considering them. Many self-evident propositions are like this. But when I consider some self-evident propositions, such as that if there never have been siblings, then there never have been first cousins, it may or may not take me a moment to see their truth. Still, when one does see such a truth, the belief one forms will (at least normally) be epistemically immediate, not inferential. So, this proposition and my coming to believe it may or may not be temporally

immediate. By contrast, the proposition that I am now seeing print is temporally immediate (for me) but is not self-evident. It is evident not in itself, but *through* what I see.

- 6 Kant's most detailed presentation of his views on these matters is in his *Critique of Pure Reason* (first published in 1781), but a short presentation is provided in the Preamble to his *Prolegomena to Any Future Metaphysics* (1783). Kant's conception of the analytic is quite reminiscent of Aquinas's idea that the self-evident has "its predicate contained in the notion of the subject" (*Summa Theologiae*, Question 94, Article 2).
- 7 There has long been controversy about whether such thought is possible without using language, or at least having a language. Donald Davidson is among those to argue for a strong dependence of thought on language. See, for example, his *Inquiries into Truth and Interpretation* (Oxford: Oxford University Press, 1984). Relevant critical discussion of Davidson is provided by Ruth Barcan Marcus in 'Some Revisionary Puzzles about Belief and Believing', *Philosophy and Phenomenological Research*, supplement to vol. 50 (1990), 133–53, which brings out serious problems for the view that beliefs must have sentence-like objects. There is no need to take a stand on this issue for my main purposes in this book.
- 8 One way to conceive this is as follows: if the concept of *F* is part of the concept of *G*, then having the property (of) *F* is self-evidently entailed by having the property (of) *G*. I do not accept this overall conception of conceptual containment but do believe that the entailment holds (even if not self-evidently).
- 9 This is plausible if (1) the correct analysis of a key concept in an analytic proposition, say that of a vixen, is discoverable, without reliance on anything beyond understanding that concept, by anyone with an (adequate) understanding of the proposition, and (2) given a correct analysis of that concept, the truth of the analytic proposition is appropriately evident. However, some analytic propositions are not understandable in this way; some might be *provable* only by a lengthy process from one that is (a notion discussed on page 113). Further, it is by no means clear that every analytic proposition is self-evident in the very common sense that implies a fairly high degree of obviousness. If, as seems plausible, the self-evidence of a proposition simply implies that *some* kind of adequate understanding is sufficient for justification for believing it, then we might plausibly distinguish between the immediately and the mediately self-evident and allow that the latter propositions may be understandable (to normal persons) only on the basis of considerable reflection. Cf. Thomas Aquinas's view (which Kant might have known) that:

Any proposition is said to be self-evident in itself, if its predicate is contained in the notion of its subject . . . *Man is a rational being*, is, in its very nature, self-evident, since he who says *man* says a *rational being*; and yet to one who does not know what a man

is, this proposition is not self-evident . . . some propositions are self-evident only to the wise, who understand the meaning of the terms of the propositions.

(*Summa Theologiae*, Question 94, Article 2)

- This seems to anticipate Kant's containment notion of the analytic and largely accords with the conception of the self-evident I have introduced.
- 10 There are philosophers who regard colors as subjective in a way that might seem to undermine the example here. I do not see that taking the proposition that nothing is red and green all over at once to be necessary, synthetic, and a priori entails any particular analysis of color properties, and I doubt that the example fails. If the example should depend on a mistaken realist account of color and for that reason fail, anti-realism about shape properties is less plausible, and the proposition that nothing is round and square might serve as well. For accounts of the status of color see C.L. Hardin, *Color for Philosophers, Unweaving the Rainbow* (Indianapolis: Hackett, 1988), and Edward Wilson Averill, 'The Relational Nature of Color', *Philosophical Review* 101 (1992), 551–88. For a detailed discussion of color properties, with application to the apparently synthetic a priori proposition that nothing is red and green all over at once and with a defense of the view that color properties supervene on (and so are determined by) dispositional properties of physical objects, see Colin McGinn, 'Another Look at Color', *Journal of Philosophy* 93, 2 (1996), 537–53.
  - 11 This allows that such propositions can *also* be known empirically, say through testimony, though there are restrictions (discussed in Chapter 7) on how this may occur. The characterization suggests that an a priori proposition is knowable non-inferentially even if only on the basis of considerable reflection, but the exact mode of the appropriate reflection is not something that need be settled here. A full account of this conception of the a priori would explicate the kind of possibility of knowledge in question; it is presumably not mere logical possibility in the sense that no contradiction is formally entailed by the occurrence of the relevant knowledge, but a conceptual possibility, roughly in the sense that such knowledge is provided for by the concept of the relevant kind of knowledge: the kind grounded in understanding propositions of the sort in question. My preference is to characterize the a priori in terms of self-evident propositions and leave open what kind of possibility there has to be of the sort of understanding that grounds justification for believing those propositions. For a valuable treatment of possibility and necessity arguing that such modal notions are irreducible, see Scott A. Shalkowski, 'Conventions, Cognitivism and Necessity', *American Philosophical Quarterly* 33 (1996), 375–92.
  - 12 Kant's section 2b of his Preamble to the *Prolegomena to any Future*

*Metaphysics* (trans. by Lewis White Beck, New York: Liberal Arts Press, 1950) opens with ‘The Common Principle of All Analytical Judgments is the Law of [non]Contradiction’ and almost immediately continues: “For the predicate of an affirmative analytical judgment is already contained in the concept of the subject, of which it cannot be denied without contradiction.”

- 13 There is a subtlety here that needs comment. Imagine that a self-evident axiom, *A*, self-evidently entails a theorem, *t*, which in turn self-evidently entails a second theorem, *t'*. Self-evident entailment (as opposed to entailment in general) is not *transitive*: *A* can self-evidently entail *t* and *t* can self-evidently entail *t'* without *A*'s self-evidently entailing *t'*. Here one could understand the conditional proposition that if *A* then *t'* quite adequately without thereby having justification for believing it. One might need the intermediate step, *t*, to achieve that justification, and it need not be discerned simply in adequately understanding the conditional itself. This possible limitation does not preclude there being *some* kind of understanding of that conditional and related concepts, such as a perfectly omniscient being might have, in virtue of which the proposition that if *A* then *t'* can be seen to be true. This shows that—as Aquinas saw in the quotation from him in note 9—there is a related notion—self-evidence *for a particular person* (or mind)—which must be distinguished from self-evidence in its basic, non-relativized form, making reference only to anyone's understanding. Still, even if what is self-evident for God might not be self-evident for us, some propositions are unqualifiedly self-evident. The case also shows that not every proposition *provable* by individually self-evident steps from a self-evident premise may be assumed to be a priori in the (moderately) broad sense of being self-evidently entailed by a self-evident proposition; for (as just explained) such a proposition might not be self-evidently entailed by a self-evident proposition.
- 14 In a broader usage, a falsehood can be called an a priori proposition provided it is an a priori *truth* that it is false. This less common usage raises no special problems but presents a terminological complication I ignore in the text.
- 15 There is much difference in judgment about how to classify the analytic. It might be considered a semantic concept by those who think of it as truth by virtue of the *meanings* of the relevant terms. It might be regarded as ontological by those who think such truths are basic to the structure of reality. For epistemology, the notion of the a priori is the more important of the two. For an immensely influential paper arguing that neither notion is clear see W.V. Quine, ‘Two Dogmas of Empiricism’, in his *From a Logical Point of View* (Cambridge, MA: Harvard University Press, 1953). Among the widely noted replies is H.P. Grice and P.F. Strawson, ‘In Defense of a Dogma’, *Philosophical Review*

- 55 (1956), 114–58. For more recent discussion of these issues see Gillian Russell, *Truth by Virtue of Meaning: A Defense of the Analytic/Synthetic Distinction* (Oxford: Oxford University Press, 2008).
- 16 Critical discussion of the question whether the a priori must be necessarily true is provided in my ‘Skepticism about the A Priori: Self-Evidence, Defeasibility, and *Cogito* Propositions’, in John Greco (ed.), *The Oxford Handbook of Skepticism* (Oxford: Oxford University Press, 2008).
  - 17 See Bertrand Russell, *The Problems of Philosophy* (Oxford: Oxford University Press, 1912), Chapters 8–10 (these chapters are reprinted in Huemer, *Epistemology*).
  - 18 See especially J.S. Mill, *A System of Logic* (first published in 1843), particularly Book II, Chapters 5–7. For a much more sophisticated critique of apriorism in mathematics and an empiricist account of mathematical truths, see Philip Kitcher, *Mathematical Knowledge* (Oxford: Oxford University Press, 1984).
  - 19 Granting it is at best not obvious how logical truths are knowable by any analysis that reveals containment relations, their negations can be clearly seen to entail contradictions.
  - 20 How broad this is depends on the notion of entailment used. I have in mind a notion for which the negation of a proposition entails a contradiction provided that the use of formal logic, supplemented only by (correct) definitions, renders a contradiction deducible.
  - 21 Someone might think all truth is a priori on the ground that it is true a priori that (1) God exists; (2) a certain universe specifiable in every detail is the best of all possible universes; and (3) God creates the best of these universes. Then, with sufficient intellectual power, one could (arguably) reason one’s way to any truth. Gottfried Wilhelm Leibniz (1646–1716) has been read as holding a view close to this (but there are reasons to doubt that he did, including considerations about divine freedom).
  - 22 Cf. W.D. Ross, explicating how Aristotelian *intuitive induction* can yield a priori knowledge: “We find by experience that this couple of matches and that couple make four matches . . . and by reflection on these and similar discoveries we come to see that *it is of the nature of two and two to make four*” (*The Right and the Good* [Oxford: Oxford University Press, 1930], p. 32).
  - 23 The proposition that  $1 + 1 + 1 = 3$  might be held to be more intuitive than the proposition that  $7 + 5 = 12$ . But, first, in practice we might need to rely on less intuitive or much more complicated arithmetic to get a good case for the possible falsehood of the original proposition; second, and more important, the simpler proposition that  $1 + 1 + 1 = 3$  will also do as a case of a necessary mathematical truth.
  - 24 For discussion of the status of the a priori in connection with geometry, see the Appendix to Laurence Bonjour, *In Defense of Pure Reason* (Cambridge: Cambridge University Press, 1998). That book is also of

interest for its criticism of Kant, who in Bonjour's view is less a rationalist about—and less plausible concerning—the a priori than is often thought.

- 25 For discussion of vagueness and its bearing on epistemological matters (as well as references to his own and others' earlier work on vagueness) see Timothy Williamson, *Vagueness* (London: Routledge, 1994) and *Knowledge and its Limits* (Oxford: Oxford University Press, 2000).



## 6 Reason II

Meaning, necessity, and provability

- **The conventionalist view of the truths of reason**  
*Truth by definition and truth by virtue of meaning*  
*Knowledge through definitions versus truth by definition*  
*Conventions as grounds for interpretation*
- **Some difficulties and strengths of the classical view**  
*Vagueness*  
*Meaning change and falsification*  
*The possibility of empirical necessary truth*  
*Essential and necessary truths*  
*Necessity, apriority, and provability*
- **Reason, experience, and a priori justification**  
*A priori beliefs*  
*Loose and strict senses of 'a priori justification' and 'a priori knowledge'*  
*The power of reason and the possibility of indefeasible justification*



## 6 Reason II

### Meaning, necessity, and provability

The radical empiricist critique of rationalism is neither the only kind empiricists can mount nor the only plausible source of objections to it. Another important approach to understanding the truths of reason and our justification and knowledge builds on the undeniable connections between how we use our language—specifically, on our linguistic conventions—and our knowledge of truths expressible in that language.

#### **The conventionalist view of the truths of reason**

We have seen the importance of analyses for understanding the a priori. Definitions of certain kinds may be considered linguistic counterparts of analyses. On one view, analytic truths may be better seen as definitional than as “analytical.” This idea needs examination.

#### *Truth by definition and truth by virtue of meaning*

To see how this approach goes, suppose that analytic propositions may be said to be *true by definition*. On the assumption that the truth or falsity of definitions turns on linguistic conventions, one can now make moves parallel to the classical ones that are expressed in terms of concepts. Thus, ‘vixen’ is definable as meaning (the same thing as) ‘female fox’; ‘female’ is part of this phrase; hence, by grasping a definition (even if we do not call it to mind) we can *see* how the proposition that all vixens are female is true. The predicate, ‘is female’, expresses part of the meaning of the subject, ‘vixen’, just as the concept of being female is part of the content of the concept of a vixen. Thus, according to conventionalism, by *appeal* to the definition of ‘vixen’ as having the same meaning as ‘female fox’, we can also *show* that the proposition that all vixens are female expresses an analytic truth.

The conventionalist may grant that in the case of synthetic truths of reason, for instance the truth that nothing is red and green all over at once, we cannot make the same moves. For the relevant color terms are indefinable, or in any case not definable in the needed way. But we can still speak of truth by virtue of *meaning* or at least *convention*, in the limited sense that it seems

to be a matter of the meanings of, or conventions governing, say, the terms 'red' and 'green', that if one of the terms applies to a surface at a time and place, the other does not. Why else would someone who sincerely denies that nothing is red and green all over at once seem to exhibit an inadequate understanding of at least one crucial term used in expressing that proposition?

What terms mean is a matter of convention. It depends entirely on agreement, usually tacit agreement, among the users of the relevant language, concerning the proper application of the term. We could have used 'vixen' differently; we in fact would have done so if the history of our language happened to differ in a certain way. Moreover, even now we could decide to use 'vixen' differently and proceed to do so.

The suggested account of the truths of reason—*conventionalism*—grounds them in conventions, especially definitional conventions, regarding meaning. Secondly, and related to this basic claim, it conceives our knowledge of them as based on our knowing those conventions. Since knowledge of conventions is reasonably taken to be empirical knowledge based on suitable observations of linguistic behavior, conventionalism (on this interpretation) turns out to be a kind of empiricism regarding the truths of reason, and it has been held by some philosophers in the empiricist tradition. The claim is not that these truths are *about* words, but that knowledge of them is *based* on empirical knowledge of linguistic usage.

### *Knowledge through definitions versus truth by definition*

Some of the points made by conventionalism are quite plausible. In grasping the definition of 'vixen' as meaning the same thing as 'female fox', perhaps we can see that all vixens are female; and under certain conditions, by appeal to the definition we perhaps can show that this truth holds. But do these points undercut the classical view? If the points hold, that may well be because of something *non-linguistic*: perhaps, *in* grasping the definition we understand the *concepts* involved and thereby see a containment relation between the concept of a vixen and that of being female. In this or some other way, understanding definitions might be a ladder by which we climb to an understanding of concepts.

Furthermore, as a proponent of the classical account might also note, it seems possible to grasp the relevant conceptual relations, and thereby already know the analytic truth, even if one does not know any such definition. Indeed, it might be only on the basis of the analytic truths one knows—such as that all vixens are female, and that all female foxes are vixens—that one is able to *construct* a definition of 'vixen'—with its present meaning—in the first place. The definition would reflect what is already true in virtue of how the concepts in question are related; the concepts are not themselves created by or grounded in linguistic conventions.

Contrary to conventionalism, then, the knowledge of analytic truths would be essential in one's route to the definitional knowledge, not the other

way around. Understanding the relations between the concepts expressed by the words in question would be the basis for judging the definitions of those words; it would not be through first knowing the truth of those definitions that one understands the conceptual relations or knows the analytic truth. Hence, knowledge of analytic truths apparently does not *depend* on knowledge of definitions or conventions, even if the former can sometimes be gained through the latter.

The more general important point implicit here is that conventionalism fails to give a good account of what grounds the *truth*, as distinct from our knowledge—or some of our knowledge—of analytic propositions. It is not *because* ‘vixen’ means the same thing as ‘female fox’ that all vixens are female. For, as we saw in assessing the empiricist view, this analytic truth does not depend on what ‘vixen’ means. This truth holds whether there is such a word or not. It could be expressed in some other language or by other English terms. It could be so expressed even if the word ‘vixen’ never existed.

There is another way to see limitations on what we can learn merely from definitions. Suppose that, although ‘vixen’ had always meant the same thing as ‘female fox’, *both* terms had meant something else, for example ‘wily creature’. In that case, ‘All vixens are female’ would still have expressed an analytic truth, but not the one it now does. It would have meant what we now mean by ‘All wily creatures are wily creatures’.

Moreover, although one can come to know that all vixens are female *through* understanding definitions of terms that now express this truth, one cannot know it wholly on the *basis* of the truth of those definitions. A route to a foundation is not itself a foundation.<sup>1</sup> To know that all vixens are female by virtue of knowing that, say, ‘vixen’ has the same meaning as ‘female fox’, we need a bridge between knowledge of linguistic convention and knowledge of vixens. Consider one thing such a bridge requires. We must be justified in believing a general principle something like this: that a proposition expressed by a subject–predicate sentence such as ‘All vixens are female’ is true if its predicate term—here ‘female’—expresses something contained in the concept designated by its subject term, here ‘vixen’. But this bridge principle is a good candidate for an analytic truth. If it is analytic, then, on pain of generating an infinite regress, one can know an analytic truth by knowing conventions only if one *assumes* some other analytic truth.

Moreover, to know, in the light of this bridge principle, that all vixens are female, we must take the relevant sentence, ‘All vixens are female’, to be the kind of thing the principle applies to, that is, to be a sentence with a predicate that expresses something contained in the concept designated by its subject. We are in effect using logic as well as knowledge of meaning to discern something about a particular sentence and to bring that sentence under a generalization about sentences. But how can conventionalism account for our knowledge (or justified belief) of the logical truths we thereby depend on, such as that if all sentences of a certain kind express truths, and this sentence is of that kind, then it expresses a truth?

The conventionalist cannot respond by doing the same thing all over again with this logical truth; for that would presuppose logic in the same way, and the procedure would have to be repeated. The problem would arise yet again. No finite number of steps would explain our justification, and an infinite number would not be possible for us, even if it would help. We could thus never account for knowledge of a given logical truth without presupposing knowledge of one. Since conventionalism presupposes (at least) logical truths of reason, in order even to begin to account for analytic ones, it cannot show—and provides no good reason to believe—either that every truth of reason is grounded in convention or even that all knowledge of such truths is grounded in convention.

### *Conventions as grounds for interpretation*

These criticisms should not be allowed to obscure a correct point that emerges from reflecting on conventionalism. The meaning of ‘vixen’ *is* crucial for what proposition is expressed by the sentence ‘All vixens are female’, that is, for what one is asserting when (in the normal way) one uses this sentence to make an assertion. Thus, if ‘vixen’ came to mean the same as ‘wily creature’, that sentence would express a falsehood, since there are plenty of wily males. But from the fact that change in what our terms mean can result in our saying different things in uttering the same words, nothing at all follows regarding whether *what* we say in using these words is necessarily true, or true at all. Those matters depend on what it *is* that we say.

There are, however, insights underlying conventionalism: truths of reason are associated with meanings; they can be known when meanings are adequately understood; and they can sometimes be shown through pointing out relations of meanings. Moreover, without conventions, our “words” could not be said to have meanings: strictly speaking, we would have no words and could not plausibly call anything true by virtue of meaning.

Important as these points about conventions are, they do not support the conventionalist view that the truths of reason themselves, or even our justification or knowledge regarding those a priori propositions, are *based* on what words mean or on our conventions for using them. For all that these points establish, our understanding of word meanings (including sentence meanings) is simply a route to our grasping of concepts and shows what it does about the truths of reason only because of that fact.

### **Some difficulties and strengths of the classical view**

Of the accounts just considered, then, the classical view of the truths of reason and our knowledge of them apparently stands up best. But there are other accounts and many variants on the ones discussed here. Moreover, I

have sketched only the main lines of the classical view and only some of the challenges to it. There are still other difficulties for it.

### *Vagueness*

Recall the problem of vagueness. Perhaps the concept of being red, as well as the term 'red', is vague. Is it, then, an a priori truth that nothing is red and (any shade of) orange all over? And how can we tell?

One answer is that although words are by and large vague, concepts are not, and what *is* red (i.e., what instantiates the concept of redness) is never orange even though we have no non-arbitrary way of precisely specifying the limits of colors. Thus, we might confront a sentence, say 'That painting has a patch that is at once red and orange', which we cannot assess until we see whether it implies the necessary falsehood that the patch is two different colors all over at once or, because of the vagueness of its terms, expresses (say) the possible truth that the patch has a single color that can be considered red just as appropriately as orange.

This answer is only the beginning of a solution to the problem of how to deal with vagueness and is less plausible for highly complex concepts such as that of a work of art. The more vague our terms, the harder it is to discern what propositions are expressed by sentences using those terms, and thus the harder it is to decide whether these sentences express truths of reason. None of this implies, however, that there are *no* clear cases of synthetic a priori truths. Perhaps the proposition that nothing is round and square, taken to belong to pure geometry, is an example. (There may also be examples in the moral domain, an important possibility considered in Chapter 12.)

### *Meaning change and falsification*

A related problem for the classical view emerges when we consider the close connection (which some regard as an equivalence) between what a term means and the concept it expresses. With this connection in mind, notice too that meaning can change gradually, as when we discover things about vixens a little at a time and thereby almost imperceptibly come to mean something different by 'vixen'. A point may then come at which it is unclear whether the term 'vixen' expresses the concept it now does or not and, correspondingly, whether or not what is then expressed by 'All vixens are female' is analytic.

This unclarity about what concept 'vixen' expresses need not give us reason to doubt, regarding the proposition which that sentence now expresses, that it is analytic; but it does show that it may be difficult to decide whether or not an utterance or sentence we have before us expresses an analytic proposition. That difficulty may drastically limit the usefulness of the notion of the analytic in understanding philosophical and other problems.

It might be argued, moreover, that on reflection the distinction between meaning change (semantic change) of the kind illustrated and falsification

of the proposition we started with does not hold. This point is likely to be pressed by those who think that the basic epistemological standard, the fundamental standard for judging whether a belief is justified or constitutes knowledge, is what is required for an overall account of experience. This broad standard is compatible both with many versions of empiricism and with some versions of rationalism.

To understand the difference between meaning change in a sentence and falsification of what the sentence is used to assert, it is helpful to contrast two cases. Compare (1) scientists' discovering that despite appearances vixens have such significant male characteristics that they are *not* really female—an outcome the classical theory says is, on the face of it, impossible—and (2) scientists' making discoveries about vixens so startling that we come to use 'vixen' in a new sense, one such that, although scientists deny that 'vixens' in this new sense are always female, what they are thereby saying provides no reason to doubt that what we *now* mean by 'All vixens are female' is true. Is there really a clear difference between (1) and (2)—roughly, between falsification of the belief about vixens we now hold and a change in the meaning of the terms we use to express it?<sup>2</sup>

Classical theorists take (2) to be possible and tend to hold that it is only because possibilities like (2) are not clearly distinguished from (1) that (1) *seems* possible. They regard the difference between (1) and (2) as clear enough to sustain their view and tend to conclude that what may seem to be a falsification of an analytic proposition is really only a change in meaning that leads us to substitute, for an analytic truth, what looks like a proposition inconsistent with it, yet is actually compatible with it. Other philosophers think that the difference is not clear at all and that future discoveries really can weigh against what the classical view calls analytic propositions.<sup>3</sup>

It is difficult to doubt, however, that there are *some* truths of reason, such as elementary logical principles, and such simple analytic propositions as that all vixens are female, that are both a priori and necessarily true.

Whether some truths of reason are also synthetic rather than analytic is more controversial, but it looks as if some of them are. Whether, if some of them are, those synthetic truths are also invariably necessary is also very controversial. I see no good reason to deny that they are necessary, but there may be no clearly decisive argument to show this.

If synthetic truths of reason are necessary, perhaps one must simply see that this is so by reflecting on the examples. In any case, our capacity of reason, our rational *intuition*, as it is sometimes (perhaps misleadingly) called, is a source of beliefs of simple truths of reason, such as the self-evident truth that if the spruce is taller than the maple then the latter is shorter than the former. We can know the *truth* of these *intuitively*, on the basis of understanding them rather than on the basis of premises for them or perceptual experience, even if more is required to know their *status* as, say, necessary or contingent, a priori or empirical. Moreover, reason, applied in our contemplating or reflecting on certain a priori truths, can yield both situational

justification—hence justification for holding beliefs of them—and actual justified beliefs of them. Clearly, reason can also yield knowledge of them.

### *The possibility of empirical necessary truth*

It is one thing to say, with the classical view, that every a priori truth is necessary; the thesis that every necessary truth is a priori is less plausible. Consider the truth that sugar is soluble in water. Ordinarily this is thought to be a law of nature and as such something that must (of necessity) hold. Yet it is not self-evident and apparently not even broadly a priori: one could adequately understand it without thereby being justified in believing it, nor does it seem to follow self-evidently from anything self-evident. Indeed, it seems to be the kind of truth that can represent an empirical discovery. Proponents of the classical view would maintain that the necessity in question is not “logical” in the sense of absolutely precluding falsehood, but *nomic* (from the Greek *nomos*, for law), in roughly the sense characterizing laws of the natural world as opposed to every possible world or situation.

It does appear that we can clearly conceive of a lump of sugar’s failing to dissolve in water, whereas we cannot clearly conceive of something that is (in overall shape) both round and square (if this is conceivable at all). But perhaps once the idea of solubility in water is properly qualified (in ways sketched in Chapter 12), there may no longer seem to be any more than a difference of degree between the two cases. I doubt that the difference is only one of degree, but let us leave the matter open and proceed to cases that pose a greater challenge to the classical view.

The truth that gold is malleable is arguably more basic to what gold is than solubility in water is to what sugar is. Is it even possible for something to be gold without being malleable? Compare the question whether a vixen could turn out to be male. This also seems impossible, but one difference is that whereas there are good ways of identifying specimens of gold without selecting them partly on the basis of malleability, there are no comparably good ways of identifying vixens without selecting them partly on the basis of being female. Still, even classical theorists grant that taking the proposition that gold is malleable to be necessary does not self-evidently commit one to considering it analytic. Critics of the classical view will maintain that it is not obvious that a specimen of gold could turn out to lack malleability, yet it is equally far from obvious that adequately understanding the proposition that gold is malleable is sufficient to justify it.

If we move to a theoretical identification statement, such as that water is  $H_2O$ , it seems even less likely that we have a proposition that is contingent rather than absolutely necessary, yet it also appears that the proposition is not a priori. The basis of our knowledge of it is confirmed scientific theory, not understanding. To be sure, there is “heavy water,” but its existence bears on the kind of hydrogen atom, not on whether water of the everyday kind

is necessarily  $H_2O$ . In any case, a different kind of example also strongly supports this conclusion that some necessary truths are empirical. This time we turn to the domain of biology.

### *Essential and necessary truths*

As the identity of human beings is normally understood, who they are is essentially tied to their parents. Is it possible that *I* might have had (biologically) different parents? Surely anyone otherwise like me but born of different parents is only a fortuitously identical “twin.” Here, then, is an empirical proposition (that I am the son of R and E) which is apparently necessary.

Notice, however, that the proposition that I have the parents I do is singular and existential, implying the existence of the particular thing it concerns (me), whereas the clear cases of necessary truth we have considered are all general and non-existential. To say that nothing is both round and square, for instance, does not entail that there is anything round or square: it says roughly that anything which is round is non-square (and vice versa), and it would be true even if all the round and square things in the universe had been destroyed (and presumably even if there never had been any except perhaps in the mind of someone contemplating creating them).

What a proponent of the classical view might say of the parentage case is that the proposition that I have the parents I do is an *essential truth*—one attributing to a thing a property absolutely essential to it, roughly in the sense that it could not exist without it—but not a necessary truth. The idea is roughly this: a necessary truth holds in any possible world or situation; an essential truth holds in, but only in, those possible worlds or situations in which what it is about exists.<sup>4</sup>

One trouble with this view is that even in a world without water, we could speak of water and  $H_2O$  as we can of what is round or square. Perhaps the best the classical view can do here is, first, to distinguish between two kinds of necessary truth, those applicable to entities that must exist, such as (arguably) numbers, and those applicable to entities that need not exist, and second, to argue that the former truths are a priori. The idea might be that necessary truths are grounded in the nature of things, and that the nature of the kinds of things that must exist is knowable through the use of reason. The nature of water must be discovered by scientific inquiry; that of the abstract property of roundness is apparent to adequate reflection.

The idea that necessary truths are grounded in the nature of (the relevant) things has some plausibility. At best, however, it does not in any obvious way apply to purely formal necessary truths, such as that if some *As* are *Bs*, then some *Bs* are *As*, where *A* and *B* are variables and do not stand for anything in particular (they figure in indicating the form of the truth in question but provide no content).



### *Necessity, apriority, and provability*

There is, moreover, a further objection to extending the idea to imply the apriority of all necessary truths. A theorem (in one sense of the term) might follow from a necessarily true proposition and thereby be a necessary truth—as what follows from a necessary truth is itself necessarily true—yet not be a priori because there is no way to know it simply through adequately understanding it or through adequately understanding its entailment by self-evident steps from something that is self-evident. We must not simply assume that every such theorem is self-evidently entailed by a *self-evident* proposition, or that some proof of it must proceed by *self-evident* steps from a self-evident proposition. This assumption is far from obvious and not self-evident, and the classical view must establish it by argument. It is not clear that a cogent one can be found.

It should be stressed, however, that although a provable proposition need not be self-evident, a self-evident proposition may be provable. Self-evident propositions are knowable *without* proof, on the basis of adequately understanding them, and hence are not, as are many theorems, *premise-dependent*. But many can be proved, and some may need proof in order to be accepted by some people.<sup>5</sup>

Moreover, even apart from those points, the only possible proof by self-evident steps from a self-evident axiom might be long; this would put the theorem a long inferential distance from the self-evident axiom(s). Granted, such a theorem would still be *provable* from what is self-evident. But simply being thus provable (yet not self-evident) entails only being what I call ultimately a priori. That status is consistent with the possibility that, for finite minds, knowledge of the proposition depends on memory. The status is thus not sufficient for an uncontroversial kind of apriority.


It appears, then, that there can be necessary truths knowable only through the work of empirical investigation or of arduous mathematical proof of a kind that cannot ground what we might call strictly a priori knowledge. Those truths, to be sure, might be both provable and knowable just on the basis of a use of reason—though knowledge based on a long proof also seems to depend on memory. Not just any use of reason, however, qualifies knowledge reached through it as a priori.

From the falsity of the classical thesis that every necessary truth is a priori, it does not follow, of course, that the classical view is mistaken in positing synthetic a priori knowledge or in claiming that every a priori proposition is necessary. (See Figure 6.1 for a brief representation of the classical and revised views of the a priori.)

### **Reason, experience, and a priori justification**

Reason—conceived roughly as our mental capacity of understanding, especially in conceptual reflection or in inference—is a basic source of belief, justification, and knowledge. Like introspective consciousness and unlike

The classical view:



The revised view:

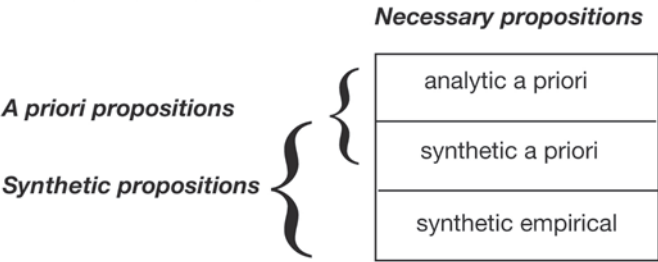


Figure 6.1 The a priori, the analytic, and the necessary.

perception and memory, it is an *active* capacity, in that we can, within limits, employ it successfully at will. I can, simply because I want to, reflect on logical and mathematical propositions. But although I can look around me just because I want to, whether I perceive anything depends on there being something there: trees and roses and books are not available to the eye in the same unfailing way that concepts and numbers are available to the mind. Through reflection on the huge range of objects of thought, we can acquire a vast amount of justified belief and significant knowledge.

To maintain that there is a priori knowledge and justification does not commit one to denying that reason has a genetic dependence on experience. Reason yields no knowledge or justified belief until experience, whether perceptual, reflective, or introspective, acquaints us with (or develops in us) concepts sufficient for grasping a priori propositions. But despite this genetic dependence of reason on experience, in one way reason may be an even firmer basis of justification and knowledge than experience. If experience is the ground from which reason grows, it is not the sole determinant of the range or power of reason. The view from the top of the tree may be more comprehensive than the view on the ground.

*A priori beliefs*

The notion of the a priori is not commonly applied to beliefs, but it should be clear from what has been said not only that it has a significant application to them but also that apriority on the part of a belief tends to indicate

some degree of justification. The following plausible *principle of justification for a priori belief* is a partial indication of the justificatory power of reason: *normally, if a rational person believes a proposition solely on the basis of (adequately) understanding it—believes it in a strictly a priori way—this belief is prima facie justified.*<sup>6</sup> In the typical cases in which this applies, the proposition, upon comprehending consideration by a rational person, will intuitively seem to the person to be true. Such an *intuitive seeming*—which for some philosophers is the primary element designated by ‘intuition’—is a source of prima facie justification. We may leave open whether this, rather than the understanding in question, is the main source of the person’s justification when the proposition in question is not self-evident. Plainly, however, the intuitive seeming presupposes at least a minimally adequate understanding of the proposition.<sup>7</sup>

There is a counterpart plausible epistemic principle—call it a *principle of knowledge for correct a priori beliefs*—to the effect that normally, if a rational person believes a true proposition in the a priori way just described, this belief constitutes knowledge. Believing in this a priori way is appropriate to (and typical for) beliefs of a priori propositions (though they may also be believed on the quite different basis of testimony), but it does not entail that the object even of a true a priori belief is a priori or a necessary truth.

It may also be true that normally, if one believes a proposition solely on the basis of one or more premises that self-evidently entail it and are themselves believed in the a priori way just described, this belief is prima facie justified. Again, such a proposition need not be a priori, but this principle is highly appropriate to what is a priori in the broad or the ultimate sense—not self-evident but either self-evidently entailed by something that is, or provable by self-evident steps from a self-evident proposition. What the principle expresses is the idea that normally self-evident entailment transmits the kind of justification that is based solely on understanding: specifically it carries that justification across a self-evident entailment. Hence, normally, if you believe a proposition on the basis of believing, with this kind of justification, a second one which self-evidently entails the first, then your belief of the first is also justified.

If these principles seem too permissive, note that we do not normally believe propositions in the strictly a priori way in question unless they are a priori and thus *can* be known on the basis of understanding them. We normally have no tendency whatever to believe, solely on the basis of understanding them, propositions indicating the state of the weather or describing the objects in our environment or the well-being or plans of others. Philosophers commonly say of such propositions that we cannot “determine a priori” (or tell or know a priori) whether they are true, and here ‘a priori’ designates an a priori way of believing rather than the status of the propositions in question. Compare how much we believe on the basis of perception, memory, and introspection; not only is this far more than is normally believed on the

basis of conceptual understanding, it is also quite different in the kind of grounding of the resulting beliefs.<sup>8</sup>

### *Loose and strict senses of 'a priori justification' and 'a priori knowledge'*

So far, I have been speaking of knowledge and justification arising from believing in a strictly a priori way. This is not necessarily a priori knowledge or a priori justification, just as not everything perceptually believed is perceptual knowledge or perceptually justified. When knowledge or justification that arises from believing in an a priori way is not strictly speaking a priori, one might still call it a priori knowledge or a priori justification in the loose sense. Let us consider justification first.

Consider the proposition that people tend to feel offended when they are insulted. This is vague, but not too vague to enable us to see that it is not an a priori truth (it seems empirically true or false, since it concerns what psychological reaction a kind of conduct in fact tends to elicit). Still, imagine someone who thinks that insulting someone self-evidently entails being offensive to the person and that feeling offended is necessarily appropriate to what is offensive and tends to occur when one is insulted. Such a person might argue that, on the basis of understanding it, we can believe the proposition that people tend to feel offended when insulted, and that we may, on this basis, be justified in believing that. If one might be so justified, then we might speak of a priori justification in the loose sense. We may also say that the belief itself is a priori in the loose sense, since it is grounded in an a priori way: if it is not grounded in the strictly a priori way (based solely on an adequate understanding of the proposition), the belief is at least held in an a priori way—it is based solely on an understanding of the proposition. Just as a perceptual belief can be justified and false (as when one first sees a straight stick half submerged in water and thinks it is bent), this belief can be also.

Another case of a priori justification in the loose sense can occur when, although one believes a proposition that *is* a priori, one believes it on the basis of an *inadequate* understanding of it. This is still believing it in an a priori way, however, as the basis of one's belief is one's understanding of the content of the proposition. But it is not believing in a strictly a priori way, as that requires adequate understanding. One might, for instance, overlook a subtlety or confuse one notion with a similar one, such as believing a proposition and being disposed to believe it. Suppose that, on the basis of my understanding of it, I believe a mathematical theorem that is a priori in the broad sense. Suppose further that this understanding, although inadequate, is not unreasonable (say because it represents a plausible though subtly misguided interpretation of the theorem). Then my belief may be justified. This is a second case of a belief held in an a priori way and exhibiting a priori justification in the loose sense. Here the proposition *is* a priori, but the justification,

though based on a reasonable understanding, is defectively grounded. In the other case of a priori justification in the loose sense, the belief is also held in an a priori way, but the proposition is not a priori.

If a belief that is a priori justified in the loose sense could constitute knowledge, we might speak of a priori knowledge in the loose sense. But as both our examples of such justification exhibit a *defective* (though reasonable) understanding in the basis of the justification, they are not plausibly considered instances of knowledge. Beliefs resting on a basis embodying conceptual error are not plausibly taken to constitute knowledge, even if the conceptual error is justified.

Suppose, however, that I believe a mathematical theorem on the twofold basis of a self-evident axiom (which I adequately understand) and the justified true belief that the theorem is entailed by the axiom (we may assume the second belief to be grounded wholly in my mathematical knowledge and understanding). Suppose further that the theorem is entailed, but not *self-evidently* entailed or self-evident.<sup>9</sup> It is not self-evidently entailed because adequately understanding the conditional proposition that if the axiom holds then the theorem does is not sufficient to justify believing this conditional. To see the truth of this conditional proposition, I must note several intermediate steps from the axiom to the theorem, so that I do not see its truth (or the entailment it expresses) on the basis of adequately understanding the proposition. Still, the entailment is provable, and by proving it I may know the theorem. This is surely a broadly a priori way of knowing it, and the proposition itself is, in my terminology, ultimately a priori. Correspondingly, we may speak of a priori knowledge in the loose sense here. The knowledge is not a priori in the *strict* sense because the theorem is not a priori, even in the indirect sense. By valid deduction, I can prove it using the a priori procedures illustrated, but such provability of a proposition is not sufficient for its being self-evident or even knowable a priori in the strict sense of that phrase.

By contrast, a priori knowledge in the strict sense is not only more than true belief held in a strictly a priori way; it is also more than knowledge of an a priori proposition. I could know a simple logical truth on the basis of testimony, even if it *can* be known on the basis of understanding alone. This would be knowledge of an a priori proposition that is not even a priori knowledge in the loose sense. Its grounding (wholly) in testimony does not prevent its being knowledge, but testimonial grounding of a belief does preclude its constituting a priori knowledge of any sort. Again, the analogy to perception is helpful. Just as perceptual knowledge is knowledge based on perception and thus more than knowledge about a perceptible, a priori knowledge is knowledge based on understanding and thus more than knowledge of an a priori proposition.

To achieve a more specific characterization of a priori knowledge we do well to begin with a crucial constituent of it—a *priori justification*. In the strict sense (the sense that mainly concerns us), this is justification based

directly or indirectly on understanding a self-evident proposition (the justification will be only situational if the person in question does not believe the proposition). A priori justification (in the strict sense) thus divides into two kinds, depending on whether it is directly or indirectly based on understanding some self-evident proposition. (1) A priori justification for believing a proposition is based *directly* on such understanding when the justification depends only on understanding that proposition itself. This is a priori justification in the strict and narrow sense. (2) A priori justification for believing a proposition is based *indirectly* on such understanding when the justification depends on *also* understanding a self-evident entailment of that proposition by some self-evident proposition. This is a priori justification in the strict but broad sense.<sup>10</sup>

If this outline is correct then *a priori knowledge*, in the strict sense, might be plausibly taken to be knowledge that is based, directly or indirectly, in the way just indicated, on understanding one or more self-evident propositions. There is, then, in addition to a division between a priori justification and a priori knowledge in the strict and loose senses, a division between direct and indirect (non-inferential and inferential) a priori justification, and direct and indirect a priori knowledge, in both senses.<sup>11</sup> (Figure 6.2 represents the four dimensions of the a priori we have been exploring.)

### *The power of reason and the possibility of indefeasible justification*

We have seen that, and perhaps to some extent how, the justificatory and epistemic power of reason enables it to ground a priori knowledge and a priori justified beliefs of a priori propositions. We have also seen its power to provide such knowledge and justification, in loose senses of ‘a priori knowledge’ and ‘a priori justification’, for propositions that are not a priori but invite belief on the basis of their conceptual content. These senses are especially appropriate for propositions that are provable from what is a priori. Is the power of reason such that it provides for something that even introspective experience apparently does not—indefeasible justification? It will help to focus on a concrete example.

There may be truths of reason that are so simple and luminously self-evident that they *cannot* be unjustifiably believed, at least at a time when one comprehendingly considers them. Could one comprehendingly consider, yet unjustifiably believe, that if Shakespeare is identical with the author of *Hamlet* then the author of *Hamlet* is identical with Shakespeare? This is doubtful. One could perhaps believe it partly on the basis of a bad argument; if one did, there would be something unjustified in the *way* one believes it. But if one believes it, one has some understanding of it, and if one understands something this simple to the extent required for believing it, it is at best difficult to see how one could fail to have an understanding of it adequate to yield

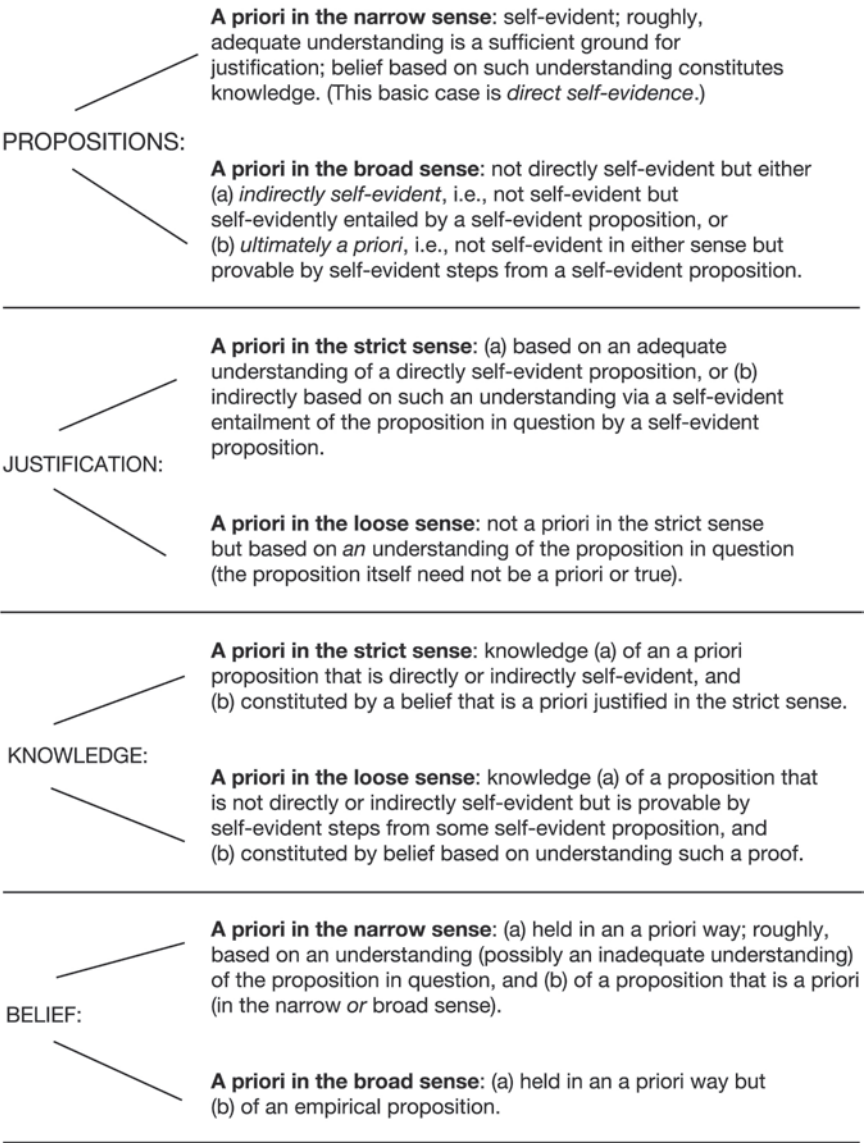


Figure 6.2 Outline of a four-dimensional conception of the a priori.

justified belief of it, at least at a time when one comprehendingly considers it. Perhaps, then, a belief held under these conditions would be—or at least could be—indefeasibly justified.

If there are propositions like this then there can apparently be indefeasible justification: justification so secure that those possessing it cannot be



unjustified in believing the proposition in question.<sup>12</sup> But not all a priori justification (even in the strict sense) should be considered indefeasible. Justification for believing even certain logical truths can be defeated by plausible skeptical arguments.

Perhaps, moreover, not all presumptively indefeasible justification need be a priori. Consider my justification for believing that I exist, a proposition that is neither a priori nor necessary but is arguably such that I cannot unjustifiably believe it. If there is indefeasible justification, this is important in dealing with skepticism (as Chapter 13 will), but plainly such justification is not a characteristic mark of either a priori or empirical justification. If, on the other hand, there is no indefeasible justification (something I leave open here), at least our understanding of simple self-evident truths of reason gives us both very secure justification for believing those truths and, when we do believe them on the basis of adequately understanding them, knowledge of them.

In summarizing some apparently warranted conclusions regarding the truths of reason, we might focus on how much seems plausible in the classical view that the a priori is coextensive with the necessary but includes the analytic as a subcategory: that any proposition that is a priori is necessary and conversely, but not every a priori proposition is analytic. Apparently, it is true that not all propositions knowable on the basis of adequately understanding them are analytic. The classical view seems correct in its claim that not everything a priori is analytic. It seems mistaken, however, in the idea that every necessary proposition is a priori, though probably not in the plausible idea that every a priori proposition is necessary.

More positively, in addition to our having a priori knowledge of self-evident propositions, on the basis of such knowledge we may know many truths that are at least ultimately a priori: not themselves self-evident but self-evidently entailed by, or provable by self-evident steps from, some proposition that is. Many of our beliefs, most clearly certain logical and mathematical ones, are grounded in understanding of their content. Reason, then, as manifested in our capacity for understanding, is one of the basic sources of belief, justification, and knowledge; and, in a way that the other three sources we have explored do not, it enables us to know truths that hold not only in the world of our experience but also in any circumstances whatever.

## Notes

- 1 At least in his classic 'Two Dogmas of Empiricism', in his *From a Logical Point of View* (Cambridge, MA: Harvard University Press, 1961), W.V. Quine sometimes talks as if he thinks that a knowledge of synonymy (sameness of meaning) of words is necessary for any possible knowledge of analytic propositions. See, for example, section 4, on semantical rules. One important comment is that "definition turned out to be a



will-o-the-wisp, and synonymy turned out to be best understood only by dint of a prior appeal to analyticity.” In the overall context, the suggestion may be that only an independent conception of synonymy would clarify analyticity.

- 2 Cf. W.V. Quine’s remark that “truth in general depends on both language and extra-linguistic fact. The statement ‘Brutus killed Caesar’ would be false if the world had been different in certain ways, but it would also be false if the word ‘killed’ happened rather to have had the sense of ‘begat’” (“Two Dogmas”, section 4). Compare saying that the *sentence* ‘Brutus killed Caesar’ would have expressed a different, and false, proposition (which is what defenders of the classical view would likely say). Has Quine provided any reason to think that the *statement* in question—understood as the historical truth we express using the sentence—would have been false if the English word ‘killed’ had meant ‘begat’?
- 3 For a valuable discussion of the notion of the analytic in relation to the conceptual, see M. Giaquinto, ‘Non-Analytic Conceptual Knowledge’, *Mind* 105, 418 (1996), 249–68. One of his major conclusions bears on the status of such cases as the proposition that all vixens are female:

What the liberated position [Quine’s, freed of behaviorism] maintains is that any belief may be rationally rejected in the light of future findings; what it has to accommodate is that some beliefs may be rationally retained even when their customary linguistic expressions become unacceptable. These [positions] are not inconsistent.

(p. 266)

- 4 The terminology of possible worlds traces especially to Gottfried Wilhelm Leibniz and has been influentially discussed in relation to a number of the issues concerning necessity and the a priori by Saul Kripke in *Naming and Necessity* (Cambridge, MA: Harvard University Press, 1980). Kripke offers a different kind of example of empirical necessities: true identity statements formed using proper names, as in ‘Hesperus is identical with Phosphorus’ (both being names of Venus). He also argues, using the example of the standard meter stick in Paris, that an a priori truth, say that the length of the standard meter stick in Paris at time *t* is 1 meter, may not be necessary. This is a highly controversial example (more often attacked than defended), which I cannot discuss here. For detailed criticism, see Albert Casullo, ‘Kripke on the A Priori and the Necessary’, *Analysis* 37 (1977), 152–9. Casullo also usefully distinguishes knowledge of the *truth value* (truth or falsity) of a proposition from knowledge of its *modal status* (its being necessarily true or false, or contingently true or false), and argues that the classical view could be mistaken in holding that the truth value of necessary propositions is

always knowable a priori yet correct in holding that their modal status is knowable a priori.

- 5 Many philosophers have taken self-evident propositions to be unprovable, e.g. W.D. Ross (*The Right and the Good*, Chapter 2), apparently following G.E. Moore and others. A simple counter-example is the proposition that if  $p$  entails  $q$  and  $q$  entails  $r$ , then  $p$  entails  $r$ .
- 6 Two comments are needed here. First, it might be desirable to widen the characterization to allow beliefs based *at least predominantly* on understanding the proposition in question (which requires understanding the concepts figuring in the proposition); but I want to avoid here the complications that arise from considering multiple bases; thus I shall not generally qualify 'based on' and similar terms. The main points in question will hold if it is taken as equivalent to 'essentially based on'. Second, although the relevant beliefs might be thought to be *always* prima facie justified, there is at least one difficulty with this: perhaps there could be an abnormal case of a kind that prevents *any* justification from arising. This is not obviously possible, since if understanding is a sufficient basis for the belief, that might arguably carry some degree of justification. In any case, the normality formulation is significantly strong.
- 7 The view that phenomenal seemings (including perceptual as well as intuitive seemings) suffice for justification is commonly called *phenomenal conservatism*. The position is defended by, e.g., Michael Huemer in *Skepticism and the Veil of Perception* (Lanham, MD: Rowman and Littlefield, 2001). For critical discussion of the view see Matthias Steup, 'Internalist Reliabilism', *Philosophical Issues*, 14 (2004), 403–24.
- 8 The quantitative comparison may be challenged by those who think we have infinite sets of mathematical beliefs (e.g. that 2 is even, 4 is even, etc.) and beliefs based on others by trivial operations, such as forming new beliefs by adding an 'or', as when, given my belief that I am seated, I form the belief that either I am seated or I am flying to the moon. That this conception of belief is mistaken will be argued in Chapter 9, which also notes relevant literature. In any case, the contrast I am drawing here would be adequately strong even without its quantitative dimension.
- 9 As indicated in explicating self-evidence, self-evident entailment (as opposed to entailment simpliciter) is not transitive. If it were then if an axiom,  $A$ , self-evidently entailed a theorem, which self-evidently entailed another, and this held for 100 steps to theorem  $T$ , the proposition that if  $A$ , then  $T$  would have to be self-evident. But reflection on axiomatic systems shows that this is not so.
- 10 This implies that even if one justifiedly believed, and knew, an a priori proposition on the basis of a self-evident axiom, but *not* on the basis of a self-evident entailment of the former by the latter (say, by a chain of non-self-evident but valid inferences instead), the justification and knowledge would still not be a priori in the strict sense—though they might be very close to it.

- 11 Three comments are needed here. First, for one's justification to be a priori, at least in the strict sense, it must not depend (epistemically) on memory. Thus, suppose there are too many self-evident premises for me to hold in mind at the same time as I understand the proposition that my conclusion follows from them. Or, suppose there are so many self-evident steps linking a single self-evident premise to a conclusion that I cannot hold them all in mind in a way that assures understanding the ultimate entailment of that conclusion by the premise. Then *my* justification for believing this conclusion is not a priori (though I may be able to prove the conclusion). Second, and related to this, so long as there can be a mind sufficiently capacious to understand the entire set of propositions in question (the premises and the proposition that if they are true, then the conclusion is also) without dependence on memory, a priori justification for *someone's* believing the conclusion is possible. Third, as in this book generally, I regard the justification referred to as defeasible (a notion considered in this chapter and again in Chapter 11) unless otherwise specified.
- 12 It might be argued, however, that if one believed such a simple self-evident proposition *essentially* on the basis of a bad argument, one would not *justifiedly* believe it, though, by virtue of adequately understanding it, one would still *have* a justification for believing it which simply fails to serve as a sufficient ground of one's belief. I leave open whether one could believe such a proposition both fully comprehendingly and essentially on the basis of a bad argument (as opposed to one's being only influenced by such an argument).

# 7 Testimony

The social foundation of knowledge

- **The nature of testimony: formal and informal**

- **The psychology of testimony**

*The inferentialist view of testimony*

*Inferential grounds versus constraints on belief-formation*

*The direct source view of testimony*

*Testimony as a source of basic belief*

- **The epistemology of testimony**

*Knowledge and justification as products of testimony*

*Testimony and memory compared*

*The twofold epistemic dependence of testimony*

- **The indispensability of testimonial grounds**

*Conceptual versus propositional learning*

*Testimony as a primeval source of knowledge and justification*

*Non-testimonial support for testimony-based beliefs*

## 7 Testimony

### The social foundation of knowledge

If our only sources of knowledge and justified belief were perception, consciousness, memory, and reason, we would be at best impoverished. We do not even learn to speak or think without the help of others, and much of what we know depends on what they tell us. Children in their first years of life depend almost entirely on others for their knowledge of the world.

If perception, memory, consciousness, and reason are our primary individual sources of knowledge and justification, testimony from others is our primary social source of them. This is why it is a primary concern of social epistemology. The distinctive situations in which testimony yields knowledge and justification are social: in each case one or more persons convey something to one or more others. There are various kinds of testimony, however, and there are many questions about how one or another kind yields knowledge or justification.

#### **The nature of testimony: formal and informal**

The word ‘testimony’ commonly evokes images of the courtroom, where formal testimony is given. Someone sworn in testifies, offering information supposed to represent what the person knows or believes. Often such testimony recounts what was witnessed first-hand, but testimony can be an expression of what we believe about something we did not witness, such as the implications of a scientific theory or the potentials of human character.<sup>1</sup>

Formal testimony differs from the informal kind in the conditions under which it is given, but not necessarily in being more credible. Testimony of the informal kind—roughly, saying something in an apparent attempt to convey (correct) information to someone else—plays a very large role in our lives and raises the question of the importance of testimony for knowledge and justification.<sup>2</sup>

For the informal giving of information, for instance in telling someone where one was last night, ‘testimony’ is too heavy a word. We could speak of ‘informing’, but this is too narrow, both in suggesting a prepared message (as in ‘Yesterday she informed me of her plan to attend’) and in (normally)

implying that what is conveyed is true. We might regard all testimony as a kind of saying. But not all saying—even apart from what is said in fiction—is testimony. Someone who says, ‘Ah, what a magnificent tree!’ is expressing a sense of the magnificence of the tree, but not giving testimony that it is magnificent, as when an arborist cites features of shape and color in supporting a claim that the tree is magnificent and worth the high cost of pruning.

For much conveyance of information it can help to speak of *attesting*. This covers both formally testifying that something is so and simply saying, in the relevant informational way, that it is so, for instance telling someone the time. Testimony is always given to one or more persons (to oneself, perhaps, in the limiting case). It may be actual or, in some cases, hypothetical, as when a diarist describing atrocities for posterity does not know whether anyone will read the testimony. In any event, what we must understand here is the role of testimony of all these kinds—roughly, of people’s telling us things—in accounting for our knowledge and justification. I begin with the psychological question of how testimony yields belief. The psychology of testimony is both intrinsically interesting and epistemologically important.

## **The psychology of testimony**

If we start thinking about testimony by focusing on formal cases, we might conclude that as a source of belief, testimony is quite unlike perception in that testimony produces in us only inferential beliefs of what is said, whereas perception produces non-inferential beliefs about what is perceived. The idea that beliefs based on testimony arise by inference from one or more premises is probably a natural result of concentration on formal testimony. When I hear courtroom testimony, I appraise the witness, place the testimony in the context of the trial and my general knowledge, and accept what is said only if, on the basis of this broad perspective, it seems true. I do not just believe what I hear, as I may just believe that a bat flew by if I see one zigzag across the evening sky. Sometimes it is like this: given the premises that (for example) the witness seems credible and that the statement in question—say that the accused dined in a certain restaurant on New Year’s Eve—fits what I know, I may thereby come to believe this statement. Let us assess the idea that testimony-based beliefs in general arise in this inferential way.

### *The inferentialist view of testimony*

If this inferentialist picture of testimony is correct, then testimony is a significantly less direct source of belief than perception: it yields belief only through both the testimony itself and one or more premises that support the proposition attested to or the attester’s credibility. If that is so, testimony is also not as direct a source of knowledge or justification; for one would know, or be justified in believing, what is attested only if one knows, or is at

least justified in believing, one's premise(s). One could not know simply *from* testimony, but only from premises *about* it as well.

There is a different, and I think more plausible, account that can also explain the psychological role of certain background beliefs. On this account, beliefs about the credibility of the attester and beliefs pertinent to the attested proposition may play a mainly filtering role. These beliefs (among other filters) prevent our believing testimony that does not "pass," for instance because it seems insincere. But if no such difficulty strikes us, we "just believe" (non-inferentially) what is attested. These filtering beliefs (which are common but not necessarily elements in the mind of everyone who forms beliefs from testimony) are like a trapdoor that shuts only if triggered. Its normal position is open, but it stays in readiness to block what should not enter.<sup>3</sup>

The open position of our natural testimonial filter is a kind of *trust*. Trust is indeed apparently the evolutionary default position: without a significant degree of trust in others, children could not, without extraordinary luck, reach adulthood and our species would likely not have survived. The absence or laxity of filtering beliefs yields *credulity*; the presence of excessively rigorous ones yields *skepticism*. Intellectual virtue—and *epistemic responsibility* conceived as a kind of virtue—are attained when we achieve a reasonable "mean" between excessive credulity and unwarranted skepticism.

It could very well turn out that, in different circumstances, each of these accounts—the inferentialist account and the non-inferential filtering account—applies to the formation of beliefs of what we are told. The psychological possibilities here are numerous, and it should be stressed that beliefs are not the only filtering elements that can non-inferentially guide formation of testimony-based beliefs. An intuitive sense of plausibility may also serve to filter. Fortunately, we need not describe all the possible filters. For now, it is enough to see that we need not consider belief properly said to be *based on testimony* to be inferential, say grounded in a further belief that the attester has spoken plausibly.

In the case of informal testimony—the most common kind—the beliefs it produces in the hearer are surely not inferential. Certainly when trusted friends speak to us on matters we have no reason to think are beyond their competence, we normally "just believe" what they tell us. Indeed, if I am trusting of people's word, then normally, when people tell me something, my belief system stands ready to be stocked. I hesitate or draw cautionary inferences only if (for instance) a would-be new belief conflicts with one or more beliefs already in my inventory. If you look vigorous and tell me you once swam the English Channel, I may readily believe you, whereas without special evidence I would not believe someone claiming to have climbed Mount Everest without using rope. For on the basis of my relevant background beliefs about climbing, I take that feat to be impossible. I have filtering-beliefs that prevent the testimony's passing into my belief system.

### *Inferential grounds versus constraints on belief-formation*

These points about how testimony produces belief need expansion. Just as it is misleading to try to build an account of the psychology of testimony from the formal cases, it is a mistake to take a *static* view of how testimony produces belief. Our beliefs and even our belief-forming processes may change in the course of our receiving testimony. I meet someone on a plane. She tells me about a conference in which a speaker I know lost his temper. Initially, I suspend judgment about whether he did so, as the incident is of a rare kind and I do not know her. Then, as she describes the conference further, other details begin to fit together very well, and she gives information I already know, such as who was there. Soon I am listening in an accepting attitude, forming beliefs of each thing she says as fast as she proceeds. At the end, I find that I *now* believe that the speaker did in fact lose his temper.

Even at the beginning, I need not have *inferred* that I should suspend judgment on the initially unlikely statement about the speaker. Suspending judgment may be a non-inferential response to the constraints set by my independent beliefs or my sense of plausibility. Moreover, her testimony is *blocked*, but not *overridden*, by my antecedent beliefs and impressions. That is, they prevent my believing what she attests to, but do not lead me to *disbelieve* it. They do not overturn a testimonially grounded belief I formed and then gave up because of what I later came to believe, as when I discover it is inconsistent with apparent facts.

What happens is apparently this. As her narrative progresses, the constraints set by my independent beliefs relax, and, regarding each statement she makes, I form beliefs not only non-inferentially, but also even spontaneously, in the sense that any constraints that might have operated do not do so. Her statements no longer have to be tested by under the gaze of my critical scrutiny, nor are they filtered out by the more nearly automatic checking the mind routinely does when people offer information.

The most difficult thing to explain here is why, at the end, I believe the proposition on which, at the beginning, I suspended judgment. One might posit an unconscious inference, say from the general credibility of her account to the conclusion that this proposition, as an essential part of it, is true. But in what sense can an inference, as a mental process, be unconscious? This is far from clear. In any case, perhaps the cognitive *influence* of my standing beliefs, such as a newly formed belief that she is credible, need not proceed through an *inference* from them. It might be like this: even apart from my forming beliefs about her credibility, her eventually becoming, in my eyes, quite credible can in some fairly direct way produce in me a general disposition to believe her. This disposition is strengthened as she speaks with an evident credibility; and at the end it overcomes the resistance to belief which was exercised earlier by my constraining beliefs. On the subject she is addressing, I have come to *trust* her. The case shows, moreover, that trust can be retroactive as well as retrospective.



### *The direct source view of testimony*

There are still other possibilities that support the conclusion that the inferentialist view of testimony is too narrow. Perhaps people (or some of us) have a credibility scale on which attesters acquire—commonly without our conscious attention to the matter—a place that can change, also without our conscious attention. This is an interesting empirical hypothesis that I cannot pursue, but all that is crucial here is that we see how beliefs grounded in testimony—testimony-based beliefs—can be *constrained by* other beliefs without being inferentially *based on* them and how beliefs based on testimony can be formed *later* than the attestation that is their ultimate source.

Perception, too, can produce belief after it has begun or, indirectly and with the help of memory, even after it has ceased. One may look at a shape for a long time before believing that it is a tree stump and not a stroller who stopped to gaze at the night sky. This same belief could also arise much later, from vividly recalling the image a day later when one is questioned about the scene. The connection in virtue of which a belief is based on a source need not be direct or simultaneous or a result of inference from premises.

Is the analogy with perception sufficient to warrant concluding that, like perception, testimony is a *basic source of belief*, in the sense, roughly, that it can produce belief without the cooperation of another source of belief? Consider perception. If I see a tree, this can produce in me a belief that there is a tree before me without my having a potentially belief-producing experience of any other sort, such as a separate consciousness of an image of a tree.<sup>4</sup> But I cannot form a testimony-based belief unless I *hear* (or otherwise perceive) the testimony. Perception is crucial for the formation of testimony-based beliefs in a way that no other belief source is crucial for the formation of perceptual beliefs.<sup>5</sup>

Granted, perception does not produce belief without appropriate background conditions, nor does its being a basic source of belief imply that antecedent beliefs are irrelevant. Suppose I firmly believe I am hallucinating the moon. Then, even if I actually see it, I may withhold judgment on whether it is out. A basic source does not derive its *generative* power from another source, but it need not operate in complete independence of other sources or their outputs. It can yield belief without the help of another source; but it may also cooperate with other sources in producing belief, and they may suppress some of its would-be products or may undermine the justification of some of the beliefs it does produce.

### *Testimony as a source of basic belief*

Given that testimony-based beliefs are not inferential, and so need not be grounded on a belief that the attester is sincere or even on a belief that someone is speaking to one (though one must be at least disposed to believe

this), one may be puzzled by the point that testimony is not a basic source of belief. The puzzlement may arise from failing to see that perception itself is required for the formation of belief based on testimony, even if perceptual *belief* is not a requirement.

To be sure, I may have to be *disposed* to believe someone has said that the speaker lost his temper to acquire a testimony-based belief of this statement; but that seems to be only because I must have comprehendingly *perceived* this *being* said; it does not imply that I have formed the belief *that* it was said, just as perception of a sentence in a convincing article one is reading can produce belief of what it says without one's forming the belief that the sentence says that. There is no reason to think the mind must keep such semantic double books. It is my perception of what is said, typically my hearing or reading it, that is required for formation of a testimony-based belief of the proposition attested to. There is a sense in which I must know—having taken in—*what* was said; but this is a kind of *understanding* and does not require forming such specific beliefs as that Juan said that Jane is reliable. We could speak of *recognitional knowledge* here: a kind of knowing *what*. Such knowledge need not be expressed in beliefs.

There is also a positive point here. Testimony can be a source of *basic beliefs*, in the sense of *beliefs not based on one or more other beliefs*. The beliefs testimony evokes need not be based on premises at all, much less on premises grounded in another belief source. The kind of non-inferential belief that testimony typically produces, the kind I am calling testimony-based belief, can also be basic knowledge *if* it meets the conditions for non-inferential (hence non-premise-based) knowledge. It can certainly be basic *for* a person in the everyday sense of being central in the person's life.

A major epistemological point that the case of testimony shows here is that a basic belief—roughly, one basic in the order of one's beliefs, and so not premise-dependent—need not come from a basic source of belief: roughly, one basic in the order of cognitive sources and so not source-dependent. A belief that is not based on, and in that sense does not depend on, another belief may come from a source of beliefs that *does* depend on another source of them.

## The epistemology of testimony

In the light of what has emerged about how testimony produces belief, we are now in a good position to ask two further questions. How does testimony yield knowledge and justification, and does it ever yield basic knowledge or basic justification in the way perception and reflection, for instance, apparently do? The case of knowledge is in some ways easier to deal with here than that of justification, and I will start with knowledge. As with perceptual knowledge and justification, testimony-based knowledge and justification turn out to differ.

### *Knowledge and justification as products of testimony*

Testimony gives knowledge to its hearers only under certain conditions. If I do not know that the speaker at yesterday's conference lost his temper, then you cannot come to know it on the basis of my attesting to it.<sup>6</sup> This is obvious if I am mistaken and he in fact did *not* lose his temper. But suppose I make a lucky guess and am right. Then I give you correct, conjectured information which I do not know; but you are also lucky to be correct and also do not know that he lost his temper. It is a fluke that I get it right; it is even more of a fluke that you get it right, since in your case there are, in addition to the chance I have taken of making a mistake, the other liabilities you escape: of my having distorted the truth, of your having misheard me, of your adding a false detail to what you take from my testimony, and so forth.

There is a more common defect in testimony that prevents its producing knowledge in the hearer. Imagine that I do not guess at, but incautiously accept, the proposition that the speaker lost his temper, from someone I know often lies about others. Again, I lack knowledge that he lost his temper, even if this time the proposition is true; and again, you cannot know it on the basis of my testimony, which is now ill-grounded in another way.

The case with justification is quite different. Even if I am not justified in believing that the speaker lost his temper, I can be credible to you in such a way that *you* can become justified in believing this on the basis of my attesting it to you. To see this, consider the two facets of *testimonial credibility*: the *sincerity* dimension, concerning the attester's honesty, and, second, the *competence* dimension, concerning the attester's having experience or knowledge sufficient to make it at least likely that if the attester holds a belief of the proposition in question or of closely related ones, then they are true. Surely you can justifiably regard me as credible on the topic of whether the speaker lost his temper if you have good reason to believe that I am honest, possess normal acuity and memory, and was present and reasonably attentive on the occasion.

Consider a further asymmetry: I cannot give you testimony-based knowledge that something is so without having knowledge that it is so, yet I *can* give you justification for believing this without having such justification. This asymmetry is important but can mislead. In both cases my credible report (my testimony) is your basis: of your justification for believing what I attest to and (when I know it) of your knowledge of it. But whereas I transmit to you my knowledge that the speaker lost his temper (when I know this), when I do not have justification for it I do not *transmit* to you justification for believing this—I do not have that justification to transmit. Rather, the way I attest to the proposition, together with your background justification regarding me and the circumstances, gives you this justification, independently of whether I have it. This illustrates non-transmissional grounding of justification, where testimony-based knowledge is transmissionally grounded.

Thus, in normal cases in which you credibly attest to something you

know, you do not give me justification in the way you give me knowledge. Testimony-based knowledge is received by what it is natural to call *transmission* and so is dependent on whether the attester knows the truth of the proposition in question—call it *p*. By contrast, recall the speaker who attests to *p*, which the person obviously knows, in a tense way that betrays anxiety, which *p* does not mention. It is natural to say that you here gain knowledge of the anxiety *through* the testimony, whereas you would gain knowledge that *p* on the *basis* of the testimony. In the first case, the knowledge has no essential relation to what is attested to—which has no connection with anxiety—and is not testimony-based.

Testimony that *p*, then, can convey the attester's knowledge that *p*; it can *produce* in the hearer *a* justification for believing *p*, and it can even yield knowledge (whether of *p* or of some other proposition) without that knowledge being based on it. But testimony that *p* does not convey the attester's justification for believing *p*—the attester need not even have such justification. My testimony that *p*, then, does not give recipients justification in the way it gives them knowledge.

This contrast between conveying knowledge and providing justification helps to explain the original asymmetry: if I do not know that a proposition is true, my attesting to it cannot transmit to you testimony-based knowledge that it is so (I have no knowledge to give here); but even if I am not justified in believing it, my attesting to it can give you justification for believing it, through providing the main materials for your becoming justified in believing it.<sup>7</sup> One might claim that this is still not testimony-based justification, but I think it can be, in the clearest sense in which there is such a thing. To see this, let us compare testimony with memory.

### *Testimony and memory compared*

The contrast between how testimony produces knowledge and how it produces justification in the recipient is reminiscent of a contrast applicable to memory (drawn in Chapter 3). Just as we cannot know that *p* from memory unless we have *come* to know it in another way, say perceptually, we cannot know that *p* on the basis of testimony unless the attester (or someone from whom the attester comes to know it) has come to know it (at least in part) in another way; whereas we can become justified in believing *p* through memory impressions, whether or not *p* is true or known,<sup>8</sup> and we can become justified in believing *p* on the basis of testimony, whether or not the attester has true belief or knowledge of it or even justification for it.

With testimony-based knowledge, as with memorial knowledge, there must apparently be a certain kind of unbroken chain from the belief constituting that knowledge to a source of the knowledge in some other mode, such as seeing; but with testimony-based justification, as with memorial justification, what is essential is apparently a matter of the present epistemic situation of the subject or recipient, such as the contents of apparently

memorial consciousness and the content and justifiedness of background beliefs. Memory and testimony can each generate justification (though in different ways); but they are not generative with respect to knowledge: characteristically, the former preserves knowledge, the latter transmits it.<sup>9</sup>

There is another way in which justification and knowledge apparently differ in their relation to testimony. Suppose I *am* justified in believing *p*, but you have no justification of your own for believing *p* or for taking me to be credible on the topic. To vary the conference example, imagine that in passing, and without giving evidence, I say that three speakers lost their tempers, and your background information neither disconfirms nor supports this claim or my credibility in the matter. Here justification follows your lights rather than mine: my would-be contribution to justifying you in believing *p* is undermined by your lack of justification for thinking my testimony is credible or for believing *p* on some other basis. Receptivity to testimony-based justification sometimes requires already having some measure of justification: for believing the attester credible or for believing *p*, or for both.

Knowledge is different on this score: to know something through my attesting to it in expression of my own knowledge, you do not have to know that I am credible; it is quite enough that you have some reason to believe I am and no reason to doubt it. It is normally enough that you presuppose it and have no reason to doubt it. Surely you can know that it is nine o'clock on the basis of my knowing this and telling it to you, even if you simply find me a normal-seeming person with a normal-looking watch and take me to be credible.<sup>10</sup> And why indeed must you meet any more than a negative condition: not having reason to doubt my credibility? After all, we are talking about a case in which I know that it is nine o'clock, attest to this *from* my knowledge of it, and thereby produce your (true) belief that it is nine. There is, then, a kind of unbroken chain from the fact that it is nine to your true belief that it is.

A natural objection to this credible-unless-otherwise-indicated view of testimony as a ground for knowledge is that in our example one's evidence is so scanty that one would at best have only some reason to believe it is nine o'clock. But is this true? Granted, my having some reason to believe the proposition may be all I can *show* from my evidence or from what I feel certain of. Still, on the assumption that I in fact do know the time and sincerely tell it to you, it would seem that you can thereby know this proposition. That appears to hold even when you simply have no reason to doubt my credibility.<sup>11</sup>

These points suggest a *principle of testimony-based justification*: At least normally, a belief based on testimony is thereby justified (i.e., justified on the basis of the testimony) provided the believer has adequate (situational) justification for taking the attester to be credible regarding the proposition in question. There is no easy way to specify the conditions for adequacy here. What we may say is that in everyday life people who do not often find people speaking falsely to them will have adequate justification for taking testimony they receive to be credible in the absence of special reason to doubt it.

We might formulate a similar principle for knowledge. To see its content let us speak of *undefeated testimony* when testimony occurs in the absence of at least the following common and probably most characteristic defeaters, i.e., factors that preclude testimony's giving the recipient knowledge: (1) internal inconsistency in what is affirmed, as when an attester gives conflicting dates for an event; (2) confused formulation, a kind that will puzzle the recipient and tend to produce doubt about whether the attester is rightly interpreted or even has a definite belief to communicate; (3) the appearance of insincerity, as when the attester seems to be lying; (4) conflict with apparent facts evident in the situation in which the testimony is given, as when a person shoveling earth over smoking coals says there has been no campfire; and (5) conflict with what the recipient knows, justifiably believes, or is justified in believing (has justification for believing). These conditions may occur separately or together; the more of them an attestation satisfies, the more clearly defeated it is, other things being equal. Trust in testimony, I take it, frees us from drawing justificatory inferences whenever we receive testimony. But if we do not have justified trust—or at least justification for having trust—then we should not believe the attester and will not be justified in believing *p* if we do.

In this light, it is plausible to hold the *principle of testimony-based knowledge*: A belief based on undefeated testimony normally constitutes knowledge provided that the attester knows the proposition in question and the believer has no reason to doubt either this proposition or the attester's credibility regarding it.<sup>12</sup> Neither this principle nor the justification principle is unqualified, but there may be only a very few cases in which abnormal conditions prevent testimony from yielding justification or knowledge (or both) when the specified conditions are met.

### *The twofold epistemic dependence of testimony*

Whatever we say about the exact conditions under which testimony grounds knowledge or justification in its recipient, we have so far found no reason to doubt that under *many* conditions testimony is a source of both knowledge and justified belief on the part of its recipient. It has seemed so far, however, that testimony cannot be a basic source of knowledge, since one cannot know something on the basis of testimony unless the attester knows it. This is why testimony does not, as such, generate knowledge though it may be described as transmitting it.

Testimony may, of course, generate knowledge *incidentally*, as when, by attesting in a surprised tone that it is 4 a.m., I give a fellow insomniac knowledge that I am awake. This knowledge is grounded not on my testimony but on the mere *hearing* of it. That kind of knowledge could as easily have been conveyed without testimony, by rising from my chair.

Testimony, like inference, can exist in indefinitely long chains. An attester might know that *p* on the basis of a third person's testimony, and the third might know it on the basis of testimony by a fourth, rather than from a generative source such as perception. But how far can this go, with each attester

informed by a previous one? There is surely some limit or other in each situation, as opposed to an infinite regress (difficulties with infinite regresses will be pursued in Chapter 9).

That brings us to a second respect in which testimony cannot be a basic source of knowledge. Surely if no one knew anything in a non-testimonial mode, no one would know anything on the basis of testimony either. More specifically, testimony-based knowledge seems ultimately to depend on knowledge grounded in one of the other sources we have considered: perception, memory, consciousness, and reason. To enable others to know something by attesting to it, I must know it myself; and my knowledge must *ultimately* depend at least in part on someone's non-testimony-based knowledge, such as knowledge grounded in seeing that the clock says four.

One might try to reinforce this view as follows. Even if someone had previously attested to a proposition, I would have to *perceive* this and to know some supporting proposition, say, that someone had credibly said it is four. But this claim is mistaken. The required kind of perceiving does not entail forming a belief of this sort, perhaps not even the specific perceptual belief that someone said it is four. The case shows, then, only that testimony is *operationally dependent* on perception, not that it is *inferentially dependent* on perceptual *belief*. It requires perceptual raw materials, but not beliefs of premises about those materials.<sup>13</sup>

If, as seems to be the case, testimony-based knowledge and justification do not depend on premises that support the testimony-based belief—say, premises confirming the credibility of the attester—this explains how such a belief can be basic. Testimony as a source of knowledge and justification need not be basic relative to other sources of knowledge and justification in order for beliefs grounded in it to be basic in the order of beliefs.

The point that testimony-based beliefs can be basic is entirely consistent with the point (made earlier) that the *attester's* knowledge that is the ground of the hearer's (potentially basic) knowledge cannot ultimately be based on testimony. Knowledge that is directly and wholly based on testimony for the recipient cannot be *ultimately* based wholly on testimony for the giver. The first would have no "right" to transmit it to the second, just as I would have no right to give someone what I had merely borrowed from someone else, who had merely borrowed it from a third person, and so on to infinity.

The point that testimony-based beliefs can be non-inferential and in that way not dependent on premises is important. But the operational dependence of testimony has both epistemological and conceptual significance. For if one did not have perceptual *grounds* for knowledge, or at least for justified belief, that someone has attested to the proposition in question, one could not know it on the basis of the testimony. This is an epistemic dependence not paralleled in the case of perception.<sup>14</sup> It shows that even if testimony-based knowledge need not inferentially depend on having *knowledge* grounded in another mode, it does epistemically depend on having *grounds*, from another mode,



grounds *for* knowledge in that other mode. Testimony-based knowledge thus depends on—and in this sense presupposes—the availability, or one might say the potential cooperation, of another source of knowledge, even if such knowledge does not require the actual operation of that source in yielding beliefs of the premises it stands ready to supply.

The case with justification is similar on this point. I cannot acquire justification for believing something on the basis of testimony unless I have justification for believing that the testifier is credible, as well as for certain other propositions, such as that I heard the testimony correctly. This justification cannot come entirely from testimony. Suppose Jane assures me about Bert, but I am not justified in taking Jane to be credible. Juan now tells me that Jane is utterly reliable. But how can this help unless I am justified in trusting Juan? Non-testimonial grounds of justification, such as perception of Juan's conduct or a memory of his guiding me in the past, must at least tacitly play some role in justifying my believing him. This role need not, however, be inferential: they need not produce in me beliefs of premises from which I infer that he is credible; they simply give me a justification that I can appeal to in framing such premises if I need them.

It may help to describe one of my overall conclusions—that testimony is not a basic source of knowledge or justification—as reflecting a disparity between the superficially simple psychology of testimony and its even more complex epistemology. Often, when we hear people attesting to various things, we just believe these things, non-inferentially and even unreservedly. But this natural psychological process yields knowledge and justification in the recipient only when certain epistemic conditions are met. In the case of testimony-based knowledge, there must be knowledge, even if not necessarily justification, on the part of the attester, whereas in the case of testimony-based justification there must be justification, even if not knowledge, on the part of the recipient. The first requirement concerns the attester's epistemic situation with respect to the proposition attested to; the second concerns the recipient's epistemic situation with respect to the attester, or the proposition, or both.<sup>15</sup>

## The indispensability of testimonial grounds

The epistemic dependence of testimony on other sources of belief must be squared with the fact that tiny children learn from what others tell them even before they are properly said to be justified (or unjustified) in believing what they do. Consider teaching color words. After a time, the child learns that the sofa, say, is red. But the tiny child has no concept of credibility or other notions important in gaining justification from testimony and, initially, insufficient experience to be justified in believing that its adult teachers are credible. On the view developed here, however, this point is quite compatible with the child's acquiring certain kinds of *knowledge*.



### *Conceptual versus propositional learning*

The first thing to note in explaining this compatibility is that there are at least two ways to learn from testimony: one can learn (in the sense of coming to know) the content attested to, and one can learn something *shown*, but not *stated*, by the testimony itself. The first case is learning *that*, specifically, that something is so. The second is learning *of* or *about* something (and may extend to learning how to do something). A tiny child just learning the basic colors is not, primarily, learning *that* (say) the sofa is red, but, above all, becoming aware of redness as the color of the sofa. This is learning colors and may be learning at least something *about* them.

In introducing the word 'red', then, the parent is only incidentally attesting to the proposition that the sofa is red. The point is to pair the word with an instance of what it stands for, with the aim of teaching that word (or, say, what the color red is), and the child can learn the main lesson without conceptualizing the sofa as such at all (something required for propositionally believing that the sofa is red). The former case of attestation—the *propositional testimony*—commonly results in propositional knowledge; we would thus have *propositional learning*. The parental introduction of vocabulary by attestation—*demonstrative testimony*—commonly results in *conceptual learning*.

It is important to see that the success conditions for the introductory function of language apparently require that for the most part the attestations be at least approximately true. A child cannot learn 'red' unless, in teaching the child English, a goodly proportion of the samples to which 'red' is applied are in fact red.<sup>16</sup> This does not of course show that most testimony is true, but it does imply that *if* communication is occurring when testimony is given to children, then one may reasonably assume that both attester and recipient have at some point benefited from a background in which a substantial proportion of attestations of a certain sort were true. How else can children be plausibly thought to have learned the language in which the communication occurs? This in turn supports the reasonableness, in everyday communicative situations, of taking testimony to be normally credible.<sup>17</sup>

At the time concepts are initially grasped in childhood, it may not be necessary that (propositional) belief and knowledge are acquired in every case. Conditions sufficient for conceptual learning may not be *automatically* sufficient for propositional learning. Belief and knowledge are, however, normally acquired at the time that concepts are initially grasped, even if conditions for mere conceptual learning are not necessarily sufficient for propositional learning.<sup>18</sup> Testimony easily produces both together. But if it cannot produce the conceptual learning without propositional learning, it can produce the latter without the former. It can be concept-producing, belief-producing (where some of the beliefs constitute knowledge), or both. The former case seems to be the more primitive, and the conditions for its possibility should not be taken as sufficient for the possibility of the latter.

It is very difficult to say when a child begins to form beliefs, as opposed to mimicking its elders by uttering things that, in adults, would express beliefs. Let us suppose both that belief-formation comes very early in life and that many of the first beliefs formed are based on what adults tell the child is the case. The child's defenseless credulity is a precondition for learning. Must this pose a problem for the epistemology of testimony suggested here?

### *Testimony as a primeval source of knowledge and justification*

Very early in their lives we speak of babies and children as knowing things. One might object that this kind of talk is simply projective: *we* would know in their situation if we behaved in the relevant way, so why not say the child does? This is a defensible response, but suppose that at least by the time children begin to talk they do know certain things. We may surely speak of their *learning*—that the milk spills when tipped, that the stove is hot, and so on—and learning (in general) implies knowledge. At about the same time, children begin learning on the basis of testimony, say that steaming tea is hot and that when the doorbell rings, someone is outside.

If, as seems a reasonable assumption, gaining testimony-based knowledge requires only having no reason to doubt the credibility of the knowing attester, then the view proposed above encounters no difficulty. If a tiny child perhaps *can* have no reason for doubt, at least the child has none; nor need there *be* any reason, since much testimony is both undefeated and unassailably credible.

Suppose, however, that a stronger requirement must be met: that the child must have (possibly in a preconceptual way) some ground for taking the speaker to be credible, for instance a series of experiences repeatedly corresponding to what the speaker says. Perhaps we could sketch a case of having such a correlational ground that would be elementary enough to fit the rudimentary character of the child's knowledge. I doubt, however, that such a ground is required for testimony-based knowledge.

With justification, there may be greater difficulty in accounting for the case of tiny children. But the first thing to notice is that we do not use the vocabulary of justification, as compared with that of knowledge, for conceptually undeveloped creatures. For a child to be justified in believing that the sofa is red, the child would have to be capable not only of having a ground for believing this but also, correspondingly, of failing to have one and yet believing this proposition anyway, thereby being *unjustified*.

It is arguable that by the time we may properly speak of children in this two-sided way as justified and also as unjustified—which is sometimes not long after they can speak—they *do* have a sense of the track record of adults in giving them information that their experience confirms. They have learned that if parents say it is cold outside, it is; and so forth. Children do not *use* the notion of credibility; but they can comprehend related concepts, such as

those needed for understanding that Mommy is right about things and baby brother must be corrected.

With testimony-based knowledge, by contrast, not even such unselfconscious justification seems required. The conditions by which knowledge is testimonially transmitted seem not to depend on justification in the recipient in the same way as does testimony-based justification. To be sure, testimony may be epistemically defeated—prevented from giving the recipient knowledge—by justified beliefs of some proposition contrary to the one attested to. But in the absence of defeaters, the recipient acquires testimony-based knowledge.

The acquisition of testimony-based justification seems to come later than that of testimony-based knowledge. One possibility for explaining how, very early in life, children may acquire an elemental kind of justification for accepting testimony is that at a very early stage they acquire a sense that they *themselves* generally give information only when they have gotten it themselves, as when they see that it is snowing or they feel hungry. For misinformation, we commonly and sometimes sternly correct children, whereas we patiently instill habits of correct reporting. This correlational sense that children apparently develop might provide a kind of analogical justification for taking others to be providing, when they give testimony, information *they* have obtained. A related and complementary hypothesis is that children have a rudimentary understanding of others in terms of what apparently explains their observed behavior. And what would explain Mommy's saying that it is snowing outside as well as her having seen that it is?

None of this is to say just when knowledge or justification enters the scene in human development, whether through the basic sources or through testimony. These are psychological questions; a philosophical theory need only leave room for plausible answers to them. The theory given here suggests that knowledge may arise before justification, but it does not entail even that. Moreover, it has at least this much harmony with the most familiar data about human development: the more natural it is, and the less figurative it seems, to speak of growing children as acquiring knowledge and justification based on testimony, the easier it is to find some elementary way in which they can satisfy the epistemic and justificational conditions set out above, such as making discriminations that enable them to assess what they are told and gaining some sense of the track record of those around them who offer information.

To say that testimony is not a basic source of justification or knowledge is not to imply that it is any less important in normal human life than a basic source. A source of knowledge and justification can be indispensable in life even if it is not basic. It may be that no normal human being would know anything apart from dependence on receiving testimony.<sup>19</sup> If there is no innate knowledge, and if we know nothing before learning a language (something I here assume for the sake of argument but wish to leave open), then unless we could acquire linguistic competence without the help of others, they would be essential in our coming to know anything at all.

If we try to imagine what would be left if we gave up all the knowledge and beliefs we have acquired on the basis of testimony, we would be quite unable to accomplish the sorting in the first place. But even beginning the task of trying to put aside what one knows on the basis of testimony suggests that one would at best be thrust back to a primitive stage of learning. I want to pursue this idea in relation to David Hume.

### *Non-testimonial support for testimony-based beliefs*

If one ponders Hume's view of testimony as capable of grounding knowledge only on the basis of a kind of validation by other sources, one may want to know to what extent testimony-based knowledge and justification, even taken item by item, can be backed up by other kinds. For Hume, our "assurance" in any matter depending on testimony "is derived from no other principle than our observation of the veracity of human testimony, and of the usual conformity of facts to the reports of witnesses."<sup>20</sup>

Let us ask whether, for each proposition one justifiably believes (wholly) on the basis of testimony, one has a justification from other sources. Call this the *focal justification question* for testimony-based beliefs. We must immediately acknowledge a complicating factor. These other sources would include propositions one justifiably believes on the basis of memory; and although one's justification for these propositions would not depend on the testimony needing support, one's beliefs of them might have been *originally* based on testimony. Much of what is stored in memory we came to believe through what others have told us in person or in writing. Still, we might be memorially justified in holding the beliefs in question even if we have forgotten their testimonial origin. Suppose, however, that we reasonably believe—as we should—that many of these beliefs have arisen through testimony. Then if we do not generally trust testimony, that reasonable belief might reduce our justification for the beliefs in question. For all that, if what we are memorially justified in believing, together with other justified beliefs we hold on the basis of our non-testimony-based experiences, can justify believing numerous propositions we find affirmed in testimony, then perhaps we do have *some* independently grounded justification for everything we justifiably believe on the basis of testimony.

Given that memory is a basic source of justification, then, we might have memorial justification for beliefs which only *seem* to be grounded in actual past experience. In any case, many of our beliefs about conditions under which people are credible are preserved in, or at least justified by, our memories. Thus, even if I have no evidence regarding *p* I may have reason to think the attester's affirming it is some reason to believe it.

To be sure, some of the memorially justified beliefs in question would not be justified unless I had been at some point justified in believing something on the basis of testimony, as when I believe one person's testimony, remember the proposition attested to, and use it in checking on another person's testimony. There may be a kind of circularity here, since testimony plays a

role in checking on the credibility of testimony. But notice two points. First, there are *two* attestations, normally by *different* attesters. Second, it might be argued that since memory is a basic source of justification, it may yield justification that supports testimony but is not testimony-based. Even if a memorially justified belief is originally justified on the basis of testimony, it may later be justified without dependence on that initial justification.

To illustrate some of these points about justifying a belief based on testimony, take a case regarding a country I do not know first-hand. Consider a radio news program announcing an earthquake in Indonesia. I have—though I may never have articulated it—a sense of the track record of the network in question and of the geology of Indonesia, a sense of how often errors of that kind are made, and so forth. One could always say that this yields a very weak justification, especially since I rely on some beliefs acquired through testimony (though that testimony may be independent of the credibility of the network in question). Certainly such a justification is far from conclusive. But there is still no good reason to think it must be inadequate.

It is natural here to raise a further question, a *global justification question*: Could one fashion an overall justification of the entire set of the propositions one believes on the basis of testimony? There are at least two questions one could be asking here. If the reference is to all the propositions one believes conjoined together—to the “long” proposition consisting of the first *and* the second *and* the third item, etc.—then one cannot even imagine contemplating such a monstrosity, much less justifying it. If, however, the reference is to the *set* of one’s testimony-based beliefs considered in the abstract, it is still not clear how to conceive justifying it. Suppose we take it to be a matter of showing that “by and large” testimony-based beliefs are justified. If we do not allow some testimony-based beliefs to justify others and we try to suspend judgment on all such beliefs we hold (assuming such massive suspension of judgment is even possible), I do not see that this corporate global justification project would work.<sup>21</sup> Let me explain.

Whatever might be possible in principle, it is doubtful that we can always avoid relying on testimony, at least indirectly, in any actual appraisal of testimony. Even one’s sense of an attester’s track record, for instance, often depends on what one believes on the basis of testimony. Think of how one news source serves as a check on another: in each case, testimony from one source is tentatively assumed and checked against testimony from another. How, then, can we globally justify testimony if we can never rely on it in the process?

There seems, moreover, not to be any general procedure by which one can produce a global justification for the proposition that the whole set of one’s testimony-based beliefs (or even a major proportion of its elements) is justified. Fortunately, that project of global justification is not one we need attempt, and the epistemology of testimony I have sketched implies on this matter at most that justified testimony-based beliefs are, to some degree, *individually* justifiable for the believer in terms of the basic sources of belief.<sup>22</sup>

Some are thus justifiable, even if it turns out that not all of them are.

Sometimes one person can confirm another's testimony simply by observing the scene described in the testimony. Indeed, the reliability of testimony, whether on a particular occasion or in a general way, can be checked through the basic sources. This is significant, for it appears that no parallel point holds for the basic sources. For instance, I cannot check on the reliability of perception, either in a particular instance or, especially, in a general way, without appealing to that very source, as when I look at something again in better light to check my color judgment. (One can use data from one sensory mode to justify a belief arrived at in another, but this is still relying on one perception to check on the reliability of perception.) Nor can I check on the reliability of memory, say by revisiting the scenes of past experiences, without presupposing that I remember the original judgments I seek to confirm by my visitation. Similar points hold for self-consciousness and (intuitive) reason. This contrast is one reason testimony is not fully on a par with the basic sources: we can confirm or disconfirm testimony without relying on it, whereas we cannot confirm a deliverance of a basic source without relying on (at least) that very source. The contrast does nothing, however, to suggest that in human life as we know it, testimony is not essential for at least a huge amount of what we know.

Testimony is a pervasive and natural source of beliefs. Many testimony-based beliefs are justified or constitute knowledge. They may even constitute basic knowledge or basic belief, both in the sense that they are not grounded in premises and in the sense that they play a pivotal role in the life of the believer. We might thus say that testimony-based beliefs not only constitute some of our basic knowledge but also are psychologically and existentially basic.

These beliefs are, however, not unqualifiedly basic epistemically. They are basic only in the sense that they are not *inferentially* dependent on knowledge or justified belief of prior premises. They are *epistemically* dependent, in a way perceptual beliefs are not, on one's having grounds for knowledge or justification, and they are psychologically dependent on one's having some ground—such as hearing someone speak—in another, non-testimonial experiential mode. Testimony-based beliefs, then, are not premise-dependent but do depend, for their epistemic or justificational status, on the basic experiential sources of knowledge and justification considered in Chapters 1–6. As a source of knowledge and justification, testimony depends both epistemically and psychologically on these other sources. This is entirely consistent, however, with its playing an incalculably important role in the normal development of our justification and knowledge.

## Notes

- 1 For a wide-ranging, historically informative account of what constitutes testimony for and numerous epistemological problems surrounding it see C.A.J. Coady, *Testimony* (Oxford: The Clarendon Press, 1992).
- 2 Not all testimony is *verbal*, much less *oral*. Consider someone's asking

of a person who requested testimony on a crime, 'What did he say?' A perfectly good answer would be 'That Mack the Knife did it', even if this was affirmed by sadly nodding when asked whether Mack is the one who did the deed. The concept of testimony allows numerous ways of telling people things; certainly any symbolic behavior rich enough to count as affirming a proposition can serve.

- 3 Thomas Reid spoke eloquently on this topic; he said, for example, "The wise author of nature hath implanted in the human mind a propensity to rely upon human testimony before we can give a reason for doing so. This, indeed, puts our judgment almost entirely in the hands of those who are about us in the first period of life." See 'Essay on the Intellectual Powers of Man' in *Thomas Reid's Inquiry and Essays*, ed. by Ronald Beanblossom and Keith Lehrer (Indianapolis: Hackett, 1983), p. 281.
- 4 Granted, I must have (and so must memorially retain) a concept of a tree; but this merely conceptual memorial state is not a potential source of belief (which is not to deny that it can play any other kind of causal role in belief-formation).
- 5 Three points may help here. First, telepathic or otherwise strange reception of testimony may, at least for our purposes, be construed as some kind of perception. Second, granting that one cannot form perceptual beliefs without having whatever additional beliefs may be needed to possess the concepts required to understand what is believed perceptually—in my example, for instance, the concept of a star-gazing stroller—this does not imply the kind of dependence on any other belief source exhibited by that of testimony upon perception. One can perceive, though not interpret, such a stroller without having these concepts; one cannot even receive testimony, and so cannot begin to interpret or learn from it, without perceiving it. Third, supposing perception cannot occur without some manifestation in consciousness (which is itself a source of beliefs), here consciousness is an element in perception in a way that perception by an audience is plainly not an element in testimony.
- 6 You might come to know it on the basis of something *about* my testimony: perhaps, for example, I give it nervously and you know that the nervousness is an after-effect of my being shaken by the fit of temper which I have since half forgotten and might even deny. This would be a case of belief *caused* by testimony but not *based* on it (not an easy distinction to explicate; but it was illustrated in Chapter 3 in showing how a belief that a past event occurred need not be a memory belief even if caused by that event; and it will be developed further in Chapter 10). One requirement for a belief to be based on testimony is the believer's holding the proposition because it was attested to, as opposed, for example, to holding it because of *how* or from what motive it was attested to. There have, however, recently been challenges to the idea that testimony-based belief constitutes knowledge only if the attester knows the proposition in



question. Some of these are cited and briefly answered in my ‘Testimony, Credulity, and Veracity’, in J. Lackey and E. Sosa (eds.), *The Epistemology of Testimony* (Oxford: Oxford University Press, 2006).

- 7 The qualifier ‘testimony-based’ is crucial: suppose I attest, in a baritone voice, that I have a baritone voice, but do not know this because I falsely believe I have a tenor voice; then you come to know, *from* my testimony, but not on the basis of it (its content), that the proposition to which I attest is true. The same point holds for justification in place of knowledge. One might also say that you come to know it *through* my testimony in a weak sense of ‘through’ not implying that the content of what I attest is crucial. It is also possible that the content, but not *my attesting it*, is essential, as when I present an argument you know I barely understand, and you come to know its conclusion, not because I attest to it or to the premises, but on the basis of yourself realizing, by bringing to bear your background knowledge, that they are true and entail the conclusion. This would be knowledge based on the *content* of testimony, but it would not be what we call ‘testimony-based knowledge’.
- 8 I develop and defend this contrast in ‘Memorial Justification’, *Philosophical Topics* 23 (1995), 251–72. Particularly interesting from the point of view of the thesis that the attester must know that *p* are examples given by Peter Graham, ‘Conveying Information’, *Synthese* 123 (2000), 365–92 and Jennifer Lackey, ‘Testimonial Knowledge and Transmission’, *Philosophical Quarterly* 49 (1999), 471–90. I will mention just one of hers. A teacher who disbelieves the theory of evolution but teaches it conscientiously tells his students, on the basis of his correct reading of the theory and his observation of a fossil, that there were *Homo sapiens* in the place in question. Since we may suppose he is giving his students correct information on a sound basis, we may tend to conclude that testimony-based belief (theirs) can be knowledge without the attester’s knowing the proposition in question. This is an interesting case, since the hearers do have a testimony-based true belief that seems well grounded. But is it knowledge, if the teacher would have taught a false theory in the same way, had his job required it? Even if the theory itself is (an item of) “knowledge” (as some would say if it is known), *he* isn’t a reliable link in the chain from the fossil through the theory, since he neither knows it nor even believes it on the kind of ground that would protect him from error in the way the (truth-conducive) grounds of knowledge do. It isn’t that the theory he uses just happens to be true, but—from the point of view of genuine evidence—he just happens to use it. If, on the other hand, we suppose that the school would not require teaching a theory that is not well evidenced, and that the students believe something to this effect, then perhaps an essential part of their basis for believing him is that background belief. Their belief would then be bolstered by background beliefs rather than a genuinely testimony-based



one. It would be as if they had to believe something to the effect that this is what the school is teaching in order to believe what he says. Chapter 10 will discuss knowledge in a way that supports this analysis.

- 9 I leave open whether knowledge transmitted by testimony can be as *well grounded* as that of the attester (though I am inclined to think it can be, say when the attester is “absolutely” reliable, a property that in principle could perhaps belong to memory in some cases). By contrast, so far as knowledge goes, “a testimonial chain is no stronger than its weakest link,” as Alvin Plantinga puts it in *Warrant and Proper Function* (Oxford: Oxford University Press, 1993), p. 84. He is speaking of what he calls *warrant*, which he views as roughly what makes true belief knowledge; and if, as I suspect, the point holds there too, then justification contrasts with warrant on this score as it does with knowledge. It should be added that if knowledge cannot be stronger than its weakest link, it probably need not be any weaker.
- 10 If this is so, it may show something else: on the assumption that you cannot know a proposition on the basis of premises you do not also know, this case would show that your testimony-based knowledge is not inferential, since the would-be credibility premise is not known but only permissibly assumed.
- 11 One possibility raised here is that of knowledge without justification. This will be considered in some detail in Chapter 10.
- 12 These principles are formulated cautiously in several ways: for instance, they allow for abnormal circumstances to provide exceptions; they allow that the resulting justification not be strong but only “adequate” for reasonable belief; they allow, but do not entail (what I think plausible but leave open), that the testimony-based belief *always* acquires prima facie justification from the testimony; and they permit the recipient to have justification or knowledge of the proposition in question from some *other* source as well. The epistemic principle might well be broadened by specifying that the recipient has no *overall* reason for doubt, but I offer that as a suggestion without adopting it.
- 13 Here I differ from Elizabeth Fricker, who (in one place) maintained that the recipient must perceptually believe “that the speaker has made an assertion with a particular content . . . capable of being knowledge . . . I have been convinced by John McDowell’s contention that hearers’ perceptions of speakers’ utterances are . . . a case of perceptual knowledge.” See ‘The Epistemology of Testimony’, *Proceedings of the Aristotelian Society* 61 (1987), 70. The reference to McDowell is to ‘Anti-realism and the Epistemology of Understanding’, in H. Parret and J. Bouveresse (eds.), *Philosophical Subjects* (Oxford: The Clarendon Press, 1980).
- 14 John Greco (in correspondence) has raised the question why consciousness is not related to perception as perception is to testimony. The beginning of an answer may be that (1) (sensory) consciousness of a ground for *p* is a *constituent* in perception that *p*, whereas no ground for *p* is a

constituent in testimony that *p*; (2) perceptual justification and knowledge entail and depend on *having* consciousness of the perceptual object which the justification or knowledge concerns; whereas (3) testimony-based justification and knowledge do not entail or depend on perception regarding what the justification or knowledge concern—namely, the proposition attested to (or even its subject matter, with which the recipient may have no relevant experience). Testimony-based justification, moreover (though not testimony-based knowledge), also normally depends on perception (or at least on sensory experience) *separate* from that required to receive the testimony; for (on the view taken in this chapter) justification for accepting the credibility of the testimony normally requires perception (or at least sensory experience) as part of the background one needs to acquire testimony-based justification. (Note 19 indicates why the normality qualification is needed here.)

- 15 The epistemology of testimony suggested here may perhaps be more stringent than that of Thomas Reid. For an interpretation and defense of the apparently Reidian view that testimony-based beliefs need not depend even for their justification on other sources of justification see Mark Owen Webb, 'Why I Know about as Much as You: A Reply to Hardwig', *Journal of Philosophy* 90 (1993), 260–70.
- 16 Strictly, the samples need only look red, as when white objects are flooded by red light; and arguably, one could even teach 'red' by producing only hallucinations of the color.
- 17 It can be connected with arguments such as we find in Donald Davidson's work for the conclusion that most of our beliefs must be true, but it does not imply that stronger conclusion. For discussion of this and other Davidsonian hypotheses, see Coady, *Testimony*, esp. Chapter 9. Cf. Tyler Burge's view that "We are a priori *prima facie* entitled to accept something that is *prima facie* intelligible and presented as true." See his 'Content Preservation', *Philosophical Review* 102 (1993), 472. Some explication and discussion of this view is provided in my 'Testimony, Credulity, and Veracity', cited above.
- 18 It is difficult to see how one could, through testimony, produce conceptual learning without producing some belief. Could a child become acquainted with what redness is in connection with being told the sofa is red, yet not acquire a belief of some sort, for example objectually believing the sofa to be red? There is no need to settle this matter here; nor can I pursue related questions concerning conceptualization in higher animals.
- 19 One reason this point is restricted to normal human beings is that it seems possible for a human being to be created, as a full-blown adult, artificially, in which case much knowledge of abstract propositions and perhaps of other sorts, such as knowledge of the perceptible external environment in which the person is made, can occur before any testimony enters the picture. The story of Adam and Eve is a theological

version of creation at the adult stage. There are also evolutionary conceptions of how knowledge first arises in human history, but these genetic questions would take us too far from our main questions.

- 20 *An Enquiry Concerning Human Understanding*, ed. by L.A. Selby-Bigge (Oxford: Oxford University Press, 1888), section 88.
- 21 We would certainly not be able to appeal to any significant segment of scientific knowledge, for there we are heavily dependent on testimony, written and oral. A plausible case that this dependence is even greater than it seems is made by John Hardwig in 'Epistemic Dependence', *Journal of Philosophy* 82 (1985), 693–708.
- 22 For supporting considerations favoring the possibility of the local justification and opposing that of a global one, see Elizabeth Fricker, 'Telling and Trusting: Reductionism and Anti-Reductionism in the Epistemology of Testimony: C.A.J. Coady's *Testimony: A Philosophical Study*', *Mind* 104 (1995), 393–411.

## **Part Two**

The structure and growth of  
justification and knowledge



## 8 Inference and the extension of knowledge

- **The process, content, and structure of inference**  
*Two related senses of 'inference'*  
*Reasoned belief and belief for a reason*  
*Two ways beliefs may be inferential*  
*The basing relation: direct and indirect belief*
- **Inference and the growth of knowledge**  
*Confirmatory versus generative inferences*  
*Inference as a dependent source of justification and knowledge*  
*Inference as an extender of justification and knowledge*
- **Source conditions and transmission conditions for inferential knowledge and justification**  
*Deductive and inductive inference*  
*Subsumptive and analogical inference*  
*The inferential transmission of justification and knowledge*  
*Inductive transmission and probabilistic inference*  
*Some inferential transmission principles*  
*Deductive transmission of justification and knowledge*  
*Degrees and kinds of deductive transmission*
- **Memorial preservation of inferential justification and inferential knowledge**

## 8 Inference and the extension of knowledge

As I sit reading, I hear knocking. I wonder whether someone is at the door. I then hear extended, very rapid knocking. It now occurs to me that it is a pecking sound, and I conclude that there is a woodpecker nearby. This way of coming to believe something differs from the way I came to believe there was a knocking in the first place. That belief was perceptual; it arose from my hearing the knocking. My belief that there is a woodpecker nearby is not perceptual. It arises not from, for instance, seeing the bird, but from a further belief: that the rapid knocking sounds like the pecking of a woodpecker. I hear the rapid knocking, recognize its character, and come to believe that it sounds like the pecking of a woodpecker. On the basis of this belief, I naturally conclude that there is a woodpecker nearby.

Some beliefs, then, arise from other beliefs and are based on them rather than directly on the sources described in Part One: perceptual, memorial, introspective, rational, and testimonial. This occurs with abstract matters as well as with perceptibles. Take mathematical proof; on the basis of knowledge of an axiom, we may infer a theorem. An *inference*—which we may think of as a kind of *reasoning*—may also proceed by way of something general to something quite specific. Studying a speech, one may determine that it is largely copied from someone else; given the standing knowledge that so representing another's work as one's own is plagiarism, one reluctantly concludes that this speech is plagiarized.

If we had only the beliefs arising directly from perception, memory, self-consciousness, reflection, and testimony, we could not build theories to explain our experience or our own view of the world. It is largely because we can inferentially build on what we already believe, that there is no limit to the richness and complexity of the ideas and theories we can construct.

The nature of the inferential processes in which one belief is formed on the basis of other beliefs is a major question in the philosophy of mind and the psychology of cognition. The way those processes can extend justification and knowledge is a major question in epistemology. Not just any inference that begins with truth must end with truth; some inferences embody poor reasoning. We can best pursue the second, epistemological question—how

inference extends knowledge and justification—by starting with the first, concerning the nature of inferential belief.

## The process, content, and structure of inference

What sort of process is it by which my belief that there is a woodpecker arises from my belief that there is a knocking which sounds like its pecking? One clue is the naturalness of saying that on the *basis* of my belief that the knocking sounds like such a woodpecker's pecking, I *conclude* that there is a woodpecker nearby. I *infer* that there is one nearby from what I believe about the knocking: that it sounds like the pecking of a woodpecker. In inferring this, I conclude something on the basis of something else.

What I conclude—the conclusion I draw—I in some sense *derive* from something else I believe. The concluding and the beliefs are mental. But neither what I conclude, nor what I believe from which I conclude it, is mental: these things are *contents* of my beliefs, as they might be of yours. They are not properties of anyone's mind, as in some sense beliefs themselves are. Such contents of beliefs—also called *objects* of beliefs—are commonly thought to be propositions (or statements, hypotheses, or something else that can be considered to be true or to be false but is apparently not itself a mental entity).<sup>1</sup>

### Two related senses of 'inference'

There are, then, two sorts of things involved when I draw a conclusion. One is the mental process of my concluding it on the basis of one or more of my beliefs or assumptions, as when I conclude that Alberto has been bitten by a deer tick on the basis of my belief (just acquired) that he has Lyme disease and my background belief that this is caused by deer tick bites. The other element in my drawing a conclusion is the set of two or more propositions which are my conclusion and my ground for it. Call the first element the *inferential process*; it is a mental episode of reasoning. Call the second its *inferential content*; it is abstract and not a process.

The inferential content indicates what is inferred from what—a conclusion from one or more premises. Specifying the content of an inference makes it obvious that inferences drawn by different people can have the same content. My inferring that there is a woodpecker nearby is drawing the *same* inference as you would make if you inferred this from the proposition that there is knocking which sounds like that of a woodpecker. Our inferrings are two different processes, one in me and one in you. But their content is the same. Sometimes 'inference' is used for the content of an inferential process. I want to talk about inference in both of these senses: as a process and as an argumental structure consisting of propositions.

If inferring is a process corresponding to a conclusion and one or more



premises for it, should we then suppose that in drawing my inference I *said* to myself something like, “Those knocks sound like a woodpecker’s; hence, there is a woodpecker nearby”? This might apply to someone learning to recognize woodpecker knocking and reasoning self-consciously, but not to me. I do not need to concentrate on the proposition that there are those sounds, much less to say to myself something like “hence there is a woodpecker.” I quickly realize, through hearing the sounds—and knowing from memory how woodpeckers sound—that they are its sounds; and on the basis of believing this proposition, I draw my conclusion without signposting my doing so by a silent ‘hence’.

### *Reasoned belief and belief for a reason*

My *drawing* the inference is something I do; it is a kind of *reasoning*. But it is not necessarily self-conscious, as is some reasoning done in order to prove a theorem from axioms. We need not introspect or even be focally conscious of our reasoning. We may instead simply draw our conclusion when our ground for it registers in our consciousness in an appropriate way. Thus, in response to wondering what I hear, I conceive the sounds as a pecking, and I then infer that there is a woodpecker. My resulting belief that there is one is, then, arrived at by reasoning and based on it. The belief may on that ground be called a *reasoned belief*; but we should also speak of a reasoned belief when such a reasoning process becomes the same kind of basis for a *previously* held belief. Thus, a belief arising non-inferentially, say from intuition or even guesswork, can later become reasoned.

Compare this with a case in which, as I am reading on an unusually still morning, a vehicle backfires in the distance. I go on reading without thinking about the noise, though I do have the thought that someone drove by. Have I *inferred*, while reading, that someone drove by, say on the basis of believing that I heard a vehicle backfire? Surely I *need* not have. I am not like someone who must think about whether the sound had certain qualities and, only after determining that it does, concludes that a vehicle passed. Being familiar with backfires, I might simply have recognized the sound as a backfire and, on the basis of this together with my standing belief that backfires are from (driven) vehicles, automatically formed the belief that someone drove by. The former belief (at least in relation to the two taken together) expresses my reason for holding the latter, which is thus a *belief for a reason*. It is not also a reasoned belief, however, because it is not arrived at by a process of reasoning. A reasoned belief is always *held* for a reason—one expressed by the premise(s) of the reasoning—but a belief (held) for a reason need not be a reasoned belief—one that is also *arrived at* by reasoning.<sup>2</sup>

This contrast between a belief for a reason and a reasoned belief may lead to the objection that I did not even form the belief that a vehicle backfired but only automatically believed, upon hearing the noise, that a vehicle passed.

This is a possible case, and it lies at the other end of the spectrum from the case in which one cautiously forms the belief that the noise is a backfire, and then self-consciously infers that a vehicle passed. But my case is intermediate: I am neither so familiar with backfiring vehicles that I “just hear” vehicles pass when I hear those sounds, nor so unfamiliar with backfires that I must go through a process of inferring that conclusion when I hear the sounds.

Similarly, suppose we are presented with premises that obviously imply a conclusion and left to infer or, on the other hand, “just see” it, as the case may be. In Shakespeare’s *Julius Caesar* we have the lines:

Yond Cassius has a lean and hungry look.  
He thinks too much, such men are dangerous.  
(Act 1, scene ii)

If we simply see and accept the implication that Cassius is dangerous, we acquire a belief for a reason, but it is not reasoned. If we need to consider these premises and then draw the conclusion as they lead one to it, we have a reasoned belief. In both cases one belief is based on another; but this occurs in different ways, both for different people and for the same person at different times.

People differ in the background knowledge and belief they bring to their experiences, and this in turn influences how those experiences produce new beliefs in them, say directly versus inferentially. Thus, in the very same situation, one person’s inference may be another’s perception. In literary interpretation or in art criticism, for instance, what the novice must discover through drawing inferences, the professional can “just see.” A similar point may apply in moral interpretation; a sensitive may literally see an injustice.<sup>3</sup>

It is an important point in the psychology of cognition that what one person believes only inferentially another believes directly, say perceptually. Both cases may be almost instantaneous, and their difference is easily missed. It is in part the failure to distinguish the cases that apparently leads some people to think that perceptual belief as such is inferential.

In seeing the difference between reasoned belief and (non-reasoned) belief held for a reason, it may help to notice that the *contexts* of the backfire and woodpecker noises differ significantly. The backfire is a kind of noise that can make it obvious that someone is driving by, whereas the pecking, far from coinciding with a flutter of wings that clearly mark the presence of a bird, is an isolated stream of sounds in the quiet of the afternoon, and it can be associated with many sources, natural and mechanical. Certainly there is an event of *belief-formation* when I hear the bang and come to believe that someone drove by. The point is that such a belief need not be reasoned: one need not form it by drawing an inference.

The contrast between a reasoned belief and a belief for a reason must not be allowed to obscure something important that is shared by the two. In both

cases, I believe one thing on the *basis* of another thing I believe; for instance, I believe that someone drove by on the basis of believing that a vehicle backfired. In both instances, then, there is an inferential (roughly, argumental) *structure* corresponding to my beliefs. It consists of a proposition we might think of as a conclusion and at least one we might think of as a premise from which the conclusion is drawn. This similarity helps to explain why there is an inclination to regard my coming to believe that someone drove by as somehow inferential.<sup>4</sup>

### *Two ways beliefs may be inferential*

There is a way to describe our two examples that helps to remind us of both their similarities and their differences. Call my reasoned belief that there is a woodpecker nearby—considered at the time I form it—*episodically inferential*, since (at the time) it arises from a process or episode of inferring, of explicitly drawing a conclusion from something one believes. Call my belief that someone drove by (also considered at the time I form it) *structurally inferential*, because, though it is not episodically inferential, it is nonetheless, as a belief for a reason, *based on* another belief in much the *way* one belief is based on a second when the first does arise from the second by inference. The first belief's being structurally inferential implies (among other things) that my holding the second belief, the basis belief (or premise belief), is at least part of what *explains why* I hold the first (presumably in a causal sense of 'explain').<sup>5</sup>

In short, episodically inferential beliefs—which are at the time reasoned—are beliefs for a reason and hence are inferential, but not every belief for a reason is episodically inferential. Beliefs for a reason are, however, roughly equivalent to those that are structurally inferential. Reasons, one might say, can lead to inferential belief by two different paths, one requiring reasoning, the other not.<sup>6</sup>

In both instances there is an inferential structure (which is no doubt reflected in the brain) corresponding to my beliefs: I believe the conclusion *because* I believe the premise(s), even though the beliefs are related by an inferential episode in one case and by an automatic process of belief-formation in the other. In the first case, I *do* something—I infer a conclusion. In the second, something *happens* in me—a belief arises on the basis of one or more other beliefs I hold. The resulting structure is much the same. Hence, a belief that is episodically inferential at the time it is formed will become structurally inferential when it is retained after the inference is drawn *if*, as is common, it remains based on the reason expressed by the premise(s). The difference is that the two beliefs arise in different ways. Only the belief that there is a woodpecker nearby, which arises by reasoning in response to curiosity, is (at the time it is formed) episodically inferential.<sup>7</sup>

### *The basing relation: direct and indirect belief*

We can also see how a belief can be inferentially based on a second without being episodically inferential if we consider a case—a kind especially important in understanding knowledge and justification—in which one *first* believes something perceptually and then the belief comes to be based on a premise. Suppose you see someone you take to be Alasdair. You do not get a good look, but believe in any case that it is Alasdair. When a friend says that she has just met Alasdair's wife at the train station, you now believe (in part) on the *basis* of her information that you saw him. But you need not, at any point, have inferred this from her information. The testimony-based belief—or indeed any new belief you form that expresses evidence that it was Alasdair you saw—can become a structurally inferential basis of your belief without your *using* it as a premise by drawing an inference. It is like an additional pillar that is placed beneath a porch after it is built and is unobserved by anyone who stands on the porch: the pillar adds support but otherwise leaves the porch as it is. The addition of this support can justify the belief it supports. If you are already justified in holding that belief, you are now doubly so. If the belief was not previously justified, it may now be justified or even constitute knowledge.

Another way to bring out what the two kinds of inferential belief have in common is to call them *indirect*. For in each case we believe one thing on the basis of, so in a sense *through*, believing another. Indirect beliefs are *mediated* by other beliefs, whether through inference or not. I refer, of course, to particular beliefs held by specific people at particular times. People differ in their inferential patterns, and these may change over time. Like backfires, woodpecking could become so familiar that, on hearing it, one just believes (non-inferentially) that it is occurring, rather as, on seeing grass, I may just believe, perceptually, that there is grass before me. That case illustrates *perceptual basing*, as opposed to the inferential kind we have been discussing. The former has been partly explicated in relation to the many instances of it described in Chapters 1 and 2, and the next chapter will clarify both kinds further.

The effects of increased familiarity show that one person's indirect belief may be another's direct belief, just as one person's conclusion may be another's premise. Similarly, my conclusion at one time can later become a basic premise, or vice versa: a proposition I believe indirectly at one time I may believe directly at another, as when I forget the premise I originally had, but I retain the proposition in memory.

There is a wide-ranging point illustrated here that is important for epistemology, psychology, and the philosophy of mind: we cannot in general specify propositions (if there are any) which can be believed only inferentially or only non-inferentially—*intrinsically basic propositions*.<sup>8</sup> Nor can we determine whether a person's belief is inferential by considering just the

proposition believed.<sup>9</sup> To be sure, it would be *abnormal* to believe (wholly) inferentially that if some dogs are pets, then some pets are dogs—in part because one would not normally believe, on the basis of a premise, what is luminously self-evident—or to believe (by sight) non-inferentially that there are seventeen cats eating scraps of beef in the backyard, since normally one would have to arrive at this on the basis of counting.<sup>10</sup> But strange cases like these are possible.

## Inference and the growth of knowledge

The examples I have given represent one way we learn through using our senses in combination with our rational powers. Through making inferences and through forming beliefs that are reason-based but not episodically inferential, we acquire not only new beliefs, but also new justified beliefs and new knowledge. Indeed, much knowledge and a great many justified beliefs arise in this way. The woodpecker case illustrates how this process works. In a single moment I come to believe, among other things, that no one is at the door and that there is a woodpecker nearby.

This is a common pattern in life: through the joint work of perception and our rational powers, particularly our inferential capacities, we acquire new beliefs, our justification is extended, and we gain new knowledge. We also forget, cease to be justified in believing certain things when we acquire disconfirming evidence, and sometimes infer conclusions we are not entitled to infer. But let us first concentrate on the way belief, justification, and knowledge develop.

### *Confirmatory versus generative inferences*

Inference is typically a source of new beliefs. But as we have seen, it need not be. Recall the backfire, and suppose I am so familiar with such sounds that categorizing them is not necessary for me to recognize them. Then I may well *directly*—i.e., *non-inferentially*—believe that a vehicle backfired. But now imagine that, realizing firecrackers have lately been set off nearby in honor of Guy Fawkes Night, I wonder whether the sound might perhaps have been that of a firecracker. Still, I do not give up, but only reconsider, my belief that it was a backfire. I recall the sound, remember that it had a muffled, not a popping, quality, and infer from its having that quality that it was indeed a backfire. Here I infer something I already believe. It is as if you arrived at a place without noticing your route and, wanting to be sure of where you are, you consider what route you must have taken. Finding a plausible route can confirm our sense that all is well, whether or not it is the route we in fact took.

My inference, then, is not a *source* of new belief, though it does in a way *alter* my belief that the sound was a backfire; the belief now becomes inferential. This is a change not in its *content*, but in its *basis*. The inference does

not produce a new belief but instead adds to my belief system a new ground for something I already believe. The inference is *confirmatory*, but not, as in typical cases, *generative*. Like an inference drawn in doing certain logic book exercises, it is not a *belief-forming inference*; but unlike many such cases (which often concern fictitious people or places), it has a conclusion that is already believed.

### *Inference as a dependent source of justification and knowledge*

Even when inference is not generative and hence is not a source of belief, it may still be a source of both justification and knowledge. Again, suppose I know that lately there have been firecrackers exploding nearby. I now might *not* know, or be justified in believing, that there was a vehicle backfire, *until* I recall the quality of the sound, rule out its being that of a firecracker, and infer in this light that a vehicle backfired. I might thus have neither justification for believing a vehicle backfired, nor knowledge that it did, until I draw the inference. Similarly, scientists who believe a hypothesis might not come to know it until, through investigating and ruling out certain alternatives, they reason their way to it, thereby inferring it, from new premises.

On the other hand, suppose I am *unjustified* in believing that the muffled sound in question represents a backfire. My situation might be this: in my whole life I have heard only one backfire; I have, however, heard many firecrackers with that sort of sound; and my belief that this sound represents a backfire is based on testimony from someone I think is usually unreliable. Here I do not become justified, inferentially or testimonially, in believing that there was a backfire. For a crucial premise of my inference—that this kind of noise represents a backfire—is one I am unjustified in believing. The same would hold if I had been unjustified in believing my other premise: that there was a muffled sound.

Now imagine a different case, this time regarding knowledge. Suppose I *am* justified in believing my premise that the muffled sound is from a backfire, as my previous experience adequately justifies my believing this. But suppose that, through no fault of my own, I have failed to discover that some common firecrackers sound precisely the same. Then, although I am still correct in believing my conclusion—that there is a backfire—I am mistaken in believing, and so do not *know*, my *premise* that this muffled sound *represents* a vehicle backfire. For it might just as well represent a firecracker. Thus, I infer a true conclusion, but from a premise that, though I justifiedly believe it, is false. This example shows something important: that I may be justified (and even correct) in believing that there was a vehicle backfire, yet not *know* that. My would-be knowledge that there was a backfire might be said to be defeated by my false premise, though my justification for believing this is not defeated.

This last case is not typical. Perhaps more often than not, inference

on the part of rational persons is a source of beliefs that are both justified and constitute knowledge. If inference is often a source of justification and of knowledge, is it a *basic* source? Our example suggests it is not. If, for instance, I am not justified in believing my premises that there was a muffled sound, and that such a sound represents a backfire, then my inferring that there was a backfire does not yield justification for my believing this conclusion, and I do not justifiedly believe it. Apparently, my inference justifies me in believing my conclusion only if I am justified in believing its premise (or premises). And apparently (though this is a controversial point to be examined in Chapter 9) that belief is justified only if it is grounded directly or indirectly in a basic source.

### *Inference as an extender of justification and knowledge*

Points like this suggest that inference is not a basic source of justification or knowledge, but rather *transmits* and thereby *extends* them, in appropriate circumstances, from one or more premises to the conclusion inferred from them. We can extend our justification and knowledge by inference, but it appears that if we have none to start with, inference, unlike perception, can give us none. Even amply justified inferences—roughly, inferences we are amply justified in drawing *given* the assumption of true premise(s)—do not *create* justification or knowledge when, because we neither know nor have justification for our premise(s), there is none to start with.<sup>11</sup>

Our examples show two kinds of inferential enhancement of knowledge and justification. The first kind is acquisition of *new* knowledge and new justified beliefs, say the knowledge that Cassius is dangerous; the second is *increase* in our justification for believing something we already hold or a buttressing of our knowledge of it, as when we infer that a vehicle backfired from a newly discovered premise to the effect that no firecrackers exploded. The first kind of inferential enhancement—*inferential extension*—yields an increase in the *content* of what we know or are justified in believing. The second kind—*inferential strengthening*—yields an increase in the *quantity* of our justification regarding the *same* content, or in the strength of our grounds for knowledge regarding the same content.

There is a third kind of enhancement of justification and knowledge that can be a variant of either sort. Consider a belief that arises by inference from two or more *independent* sets of premises, such as evidence of Cassius's being dangerous presented at the same time by two independent observers. Here we get new content by two or more pathways. In the same way, we may also acquire more justification (or stronger grounds) for what we already believe or know, as when we believe a vehicle backfired. Thus, we might have better justification for believing (or better grounds for our knowledge of) what the witnesses jointly attest to than we have for believing that on the basis of the evidence of any one of them alone.

Moreover, our experience often leads to inferential enhancement of all



three sorts without our making any particular effort to draw inferences. For the formation of structurally inferential beliefs, and even of many episodically inferential beliefs, occurs quite frequently and often spontaneously. As a timber can be silently and unobtrusively placed beneath a porch in a way that supports it, one belief can, without our noticing it, provide support for another that is already in place, or even lead to the formation of a belief we did not previously hold.

### **Source conditions and transmission conditions for inferential knowledge and justification**

If inference is not a basic source of justification and knowledge, but transmits it, it must meet two kinds of conditions. One kind concerns the premise(s) of the inference—its foundations, so to speak—the other concerns the relation of the premise(s) to the conclusion—how well those evidential pillars support what is built on them, for instance whether or not they express strong evidence for believing it. Let us take these in turn.

First, there are *source conditions*, as our examples show: one needs justification or knowledge in the first place. To see what the second kind of condition is, suppose I do know that the muffled sound I hear represents a *vehicle* backfire and I infer that a *truck* backfired. But imagine that I really cannot tell the difference between car and truck backfires. Then I do not know, in virtue of my inference, that a truck backfired. I started with knowledge, but it was not transmitted to my belief of my conclusion, as the premise from which I inferred that conclusion did not adequately support it. There was, we might say, no evidentially adequate pillar to ground my conclusion in my premises.

There are, then, *transmission conditions*, as well as source conditions, that an inference must satisfy in order to yield knowledge of its conclusion. Chapters 1 through to 7 in effect deal with source conditions in some detail, for example with how perception yields non-inferential knowledge that can provide premises for inference. Thus, I say little here about source conditions and concentrate on transmission conditions.

### ***Deductive and inductive inference***

We can best understand transmission conditions if, as is common in discussions of logic, we divide inferences into two categories, deductive and inductive. The usual basis of this division is an interpretation of the character of the inferential structure underlying the process of inference, or at least a choice of the kind of standard appropriate for assessing that structure. We can simplify matters by calling these structures *arguments*, even though they need not represent anyone's actually arguing *for* something or *with* anyone.

In this abstract sense of 'argument', an argument is discernible even when, simply to assure myself that I was correct in believing that there was a vehicle



backfire, I inferred, after reconsidering the kind of noise I heard, that there was indeed a backfire. I relied on the argument from propositions about the character of the noise to the conclusion that a backfire occurred, even though I already believed this conclusion. We use arguments not only to produce belief but also to justify something we already believe. I did this by tapping a justified source that transmitted its justification to my belief that a vehicle backfired. A natural interpretation of the case is this: I *reasoned* from the premises that (1) the noise represented a backfire and (2) if it did represent that, then there was a vehicle backfire, to the conclusion that (3) there was a vehicle backfire.

My argument here, and hence my reasoning—from the premises of the argument to its conclusion—is (deductively) *valid*; that is, it is absolutely impossible (in a logical sense) for the premises, (1) and (2), to be true and the conclusion, (3), false. In short, the premises of a valid argument, or of valid reasoning (logically), *entail* its conclusion. It is of course not in general impossible for the *premises* of valid deductive arguments to be false, and many such arguments have premises that are false. But it is absolutely impossible that the premises be true *and* the conclusion false.<sup>12</sup>

In the most careful terminology, ‘valid’ applies only to deductive arguments and, correspondingly, to valid deductive reasoning (the kind of reasoning whose essential content *is* a valid argument). We might conceive deductive reasoning as the sort that “aims” at validity, in the sense that it is of a kind best evaluated as valid or invalid. Thus, even though the argument from hallucination (discussed in Chapter 2) is invalid, the philosophical reasoning that employs it seems meant to be valid and is appropriately assessed as deductive.

By contrast, much reasoning that is not valid is simply not meant to be deductive in the first place. Suppose, for instance, that my reasoning had run (1) the noise sounds like that of a backfire; (2) the likeliest explanation of the noise is that a vehicle backfired; probably, then, (3) a vehicle backfired. As ‘probably’ signals, I do not take my reasoning to be valid or to be deductive at all: I simply take its premises to provide some reason to believe its conclusion. Even if I had not used ‘probably’, it would be inappropriate to consider this reasoning deductive. For obviously even the likeliest explanation *need* not be true; it would thus be a mistake to regard such reasoning—or the person using it—as aiming at validity.

We *could* call such probabilistic reasoning ‘inductively valid’, meaning roughly that relative to its premises there is a high probability that its conclusion is true. High probability is usually taken to be such that it is reasonable to accept a proposition having it. But to avoid confusion I simply term reasoning of that sort ‘inductively good’ (or ‘inductively strong’).

Moreover, it is reasoning processes and not abstract structures that I call deductive or inductive. I do not take arguments, as abstract structures, to be intrinsically of either kind, though we *speak* of them as deductive or inductive so far as they seem best assessed by deductive or inductive standards.

(The intentions of those presenting them are one among many other factors determining the appropriate standards for classifying reasoning processes.)

I want to stress in passing that we should not conceive deductive and inductive reasoning as they have often been characterized. Deductive reasoning has been described as “going” from the general to the particular, say from (1) all human beings are mortal and (2) Socrates is a human being to (3) Socrates is mortal. But our deductive backfire case, embodying the valid argument from (1) and (2) to (3), is different; it is about only particular things. Even in the classical example about Socrates, one premise is particular, in the sense that it concerns a single individual.

### *Subsumptive and analogical inference*

Even those who take deductive reasoning to go from the general to the particular should recognize that the reasoning from (1) all humans have fears and (2) all who have fears are vulnerable to (3) all humans are vulnerable is deductive (and valid). Perhaps they focus on cases in which we draw a conclusion about something or someone, say Socrates, by *subsuming* the person or thing under a generalization about similar entities, say people. Call such inference *subsumptive reasoning*—or *instantial reasoning*, since Socrates is supposed to instantiate the truth that all humans are mortal. Not all deductive reasoning is subsumptive.

As for inductive reasoning, it has often been said to “go from” the particular to the general, as when one infers that everyone has fears from the enumerative premises that Abe does, Beatrice does, Carl does, Donna does, and so on. This characterization is good so far as it goes. But it does not apply to reasoning—sometimes called *abductive*—from a premise stating the likeliest explanation of a presumed fact, to the conclusion that the proposition expressing that explanation is true; for instance, from the premise that the noise is best explained by a vehicle’s backfiring, to the conclusion that the noise represents such a backfire. Nor does the characterization do justice to certain reasoning by *analogy*, such as my concluding that a plant probably has a property, say hardness, because it is much like (analogous to) another plant that has that property.

It is better, then, to think of inductive reasoning as reasoning that, first, “aims” at providing good grounds for its conclusion, but not at validity, and, second, is best evaluated in terms of the degree of probability of its conclusion relative to its premises. This conception has the further advantage of applying to all three main kinds of inductive reasoning: generalizational, explanatory (abductive), and analogical.

### *The inferential transmission of justification and knowledge*

We are now in a position to explore the conditions for transmission of justification and knowledge.

Clearly the success of transmission is partly a matter of the status of the underlying argument: the one whose premise or premises are one's basis for the belief in question. The natural thing to say initially is that justification and knowledge are transmitted in deductive inference *only if* the underlying argument is valid and, in inductive inference, *only if* the underlying argument is (inductively) good (I use 'inference' rather than 'reasoning' here because the former is preferable for the wide range of contexts we are exploring). But these principles, though probably correct, need clarification. Let us consider the cases of inductive and deductive transmission separately.

Suppose Luigi hastily infers from the propositions (1) all opera lovers appreciate *The Magic Flute* and (2) Wilhelm appreciates *The Magic Flute* that (3) Wilhelm is an opera lover. This is invalid deductive reasoning and, even with true premises that Luigi is justified in believing, it would not transmit either justification or knowledge from his beliefs of them to his belief of its conclusion. Bad reasoning cannot realize the evidential potential of good premises.

Suppose Luigi then produces the better argument from (1) all opera lovers appreciate *The Magic Flute* and (2) Wilhelm appreciates it in the way one would expect of an opera lover, to (3) Wilhelm is an opera lover. Suppose we conceive his reasoning as deductive, say because Luigi's *underlying* principle—roughly, the one by which his reasoning is actually guided—is not the expected inductive one—that if all *As* are *Bs* and *x* is a *B* of a kind that might well be expected to be an *A* then *probably* *x* is an *A*—but the false principle that if all *As* are *Bs*, and *x* is a *B* of a kind that might be expected to be an *A*, then *x* is (certainly) an *A*. Then we must also say that transmission is blocked because his reasoning is invalid. He adheres to a mistaken (deductive) logical standard and hence does not acquire a justified belief through his inference.

Apparently, then, *deductive transmission of knowledge or justification requires validity*. Specifically, the argument *underlying* an inferential belief—i.e., the argument whose premise(s) constitute(s) what that belief is inferentially based on—must be valid if knowledge or justification is to be deductively transmitted to that belief from the premise belief(s) it is based on. To be sure, I could have independent grounds, such as testimony about Wilhelm, on which I know my conclusion. But if I do not have such grounds, then I cannot come to know this conclusion *through* deductive transmission of my knowledge from premises I have for it if those premises do not entail it, and hence the argument from them to it is invalid. We cannot build anything as solid as knowledge on weak supports, even if they themselves rest on a good foundation; the structure is still defective. (The case with justification is more complicated.)

### *Inductive transmission and probabilistic inference*

The case with inductive reasoning is more complicated. For one thing, the notion of good inductive reasoning is highly vague. It might seem that we

could simply define it as reasoning with premises that render its conclusion more likely than not to be true. But this will not do, though such reasoning may be called *probable* to indicate that it has this specific merit: formally, that of having a probability greater than  $\frac{1}{2}$ . Two points are important here.

First, a probability of just over 0.50 (indicating just over a fifty-fifty chance of truth) allows that even given the truth of the premises, the falsehood of the conclusion is almost as likely as its truth (since probabilities range from 0 to 1, with 0.50 [taken to represent  $\frac{1}{2}$ ] indicating the same likelihood of truth as of falsehood on the part of the proposition in question). One would not want to describe reasoning as good when its premises give its conclusion a probability of truth of just over 50 percent.

Second, judging how good one's inductive reasoning is may require assessing one's justification for the conclusion in relation to more than the premises from which one infers it. This holds, at least, when we are viewing the reasoning as occurring in a context in which various kinds of information are *accessible* to the reasoner. Relative just to the premise that Dave has a certain kind of cancer, the probability of the conclusion that he will die of it may be 0.60, as 60 percent of its victims do; but relative to his youth, vigor, and good medical treatment, the probability of his death from it may be 0.08. Thus, one's inductive reasoning from the premise that he has the particular cancer to the conclusion that he will die of it ignores relevant evidence and is not good, even though the conclusion has a probability of more than 0.50 relative to its premise.

Given the way in which such inductive reasoning can be negatively affected by new information it is sometimes called *defeasible reasoning*. By contrast, the entailment relation between the premise(s) of valid (deductive) reasoning and its conclusion is invariant: no additional information affects the entailment. In probabilistic language, the probability of its conclusion given its premises must be 1 even if further information is added, whereas adding information to the premise(s) of inductive reasoning can render its conclusion improbable relative to the new set of premises, hence defeat the original reasoning, as shown by the cancer case.<sup>13</sup>

Suppose we assume for a moment that good inductive reasoning has premises taking account of *all* the relevant evidence. May we then conclude that justification and knowledge are inductively transmitted only by inductive reasoning that is good in this comprehensive sense? This view is too strong. For often some of the relevant evidence is not *needed* for such inductive justification of one's belief because one's premises already contain sufficient evidence. Evidence may be relevant to a belief without being needed for its justification, as when testimony from a twelfth witness who agrees with the rest is unnecessary though perfectly relevant. The point is important; for even if we can understand the notion of *all* the relevant evidence, we at best rarely have all the evidence relevant to a belief and we may not need it all if we do.

Is good inductive reasoning simply the kind of inductive reasoning that is

*sufficient* to transmit justification? This is a promising characterization for single pieces of inductive reasoning, those using a set of premises directly for one conclusion. But inductive reasoning can occur in *chains*, with the conclusion of the first piece of reasoning serving as a premise in the second piece, and the conclusion of that serving as a premise in the third, and so on. Unfortunately, in an inductive chain that is extended, over time, through many inferences, justification may not be transmitted from the conclusion of the first to the conclusion of the last, even if each piece of reasoning has premises giving high probability to *its* conclusion. To see why, notice first that the *degree* of justification inductively transmitted from one's premises to one's conclusion may drop, even if nothing new enters the picture, such as someone's challenging one's conclusion the moment one draws it.

If the degree of probability repeatedly drops, the degree of justification may drop drastically. To see this, notice that even if one starts with excellent justification for one's premises, if they give a probability of only, say, 0.75 to one's conclusion, one will have much weaker justification for the conclusion than for the premises, if they are one's *only* basis for it. (I am assuming, somewhat artificially, that justification admits of degrees in the numerical way probability does.) Roughly, one should take the chance that the conclusion is true to be only 75 percent of the chance that the premises are true.

Suppose that I know that Tom said that the weather forecaster predicted rain. If the chance that Tom (who is biased by optimism and may have misheard the forecast) is right is only 75 percent and the chance that the forecaster's prediction is right is, say, 60 percent, then my warrant for believing it will rain is presumably just 75 percent of 60, i.e., 45 percent. (The idea is that the probability that the forecast was even made is only 75 percent, and we would then have a 60 percent probability of rain; the multiplication takes account of both probabilities.) Such chains of inference can be indefinitely long, as when I must rely on hearsay I take to be less than 100 percent likely to be right for my belief that Tom said the forecaster predicted rain. This allows for the occurrence of even more reduction of one's justification for believing one's conclusion.

These points should make it apparent how it is possible for good inductive reasoning, carried out through a series of inferences, to fail to transmit justification from its initial premises to its final conclusion. Even if the probability that the initial premises give to the first conclusion is 0.9, if one went on inferring further conclusions, each being a premise for the next conclusion, then even with the same degree of probability in each case, one could eventually infer a conclusion for which one has less justification than 0.5. With each case, the likelihood that one's conclusion is true would be 10 percent less than (90 percent of) the likelihood of the truth of one's previous conclusion, which is serving as one's premise.

In some respects, *knowledge* differs from justification in relation to transmission conditions. Since knowledge does not admit of degrees (at least not in the way justification does), it might be transmitted across an inductive

inference without diminution in degree even if such transmission does imply some reduction in one's justification (other things being equal).<sup>14</sup> If, for instance, you know that the weather is bad and you inductively infer that Jane, who is driving, will be late, presumably you could know the latter proposition on the basis of the former even though there is *some* chance that she left early and compensated for the weather: the probability of this is above zero. Your *grounds* for your conclusion may not be as good as your grounds for your premise, which may render the conclusion only very probable, rather than entailing it; but you may still unqualifiedly know your conclusion. This knowledge may not be as good, say as securely grounded, as your knowledge of the premises; for instance, it might not be as nearly certain. But it can still be knowledge.

Although there are *kinds* of knowledge, then, apparently a belief either constitutes knowledge or falls short of that, as opposed to constituting knowledge to a degree. This point applies to testimonial chains as well as to inductive inference. On the view taken in Chapter 7, testimony-based knowledge is non-inferential; hence, even the tenth (or indeed *n*th) recipient of testimony that *p* can know that *p* on the basis of testimony if each previous attester did.

It can happen, however, that knowledge is not transmitted even across an inductive inference whose premises give its conclusion extremely high probability. For example, you might know that you hold just one out of a million coupons in a fair lottery, which will have one winner. You may inductively infer, with very high probability, 0.999999, that you will lose, as 999,999 of the million coupons will lose. But surely you do not *know* you will lose—though, to be sure we can imagine contexts in which one would not be opposed for saying this.<sup>15</sup> You might be lucky. Moreover, you have as good a chance to be lucky as any other holder of a single coupon—including the possessor of the winning one. If we said, on this basis alone, that you know you will lose, how could we avoid having to say it of everyone else—in which case we would be wrong, since *someone* wins in this kind of fair lottery? Your knowledge of your premises, then, is not inductively transmitted to your conclusion. Granted, if we change the example so that you *deduce* the qualified statement that the probability of your losing is 0.999999, you may know *that*. But that is a very different conclusion.<sup>16</sup>

### *Some inferential transmission principles*

We have seen some important points. Inference transmits justification and knowledge; it is not a basic source of them. It can generate them only derivatively, by transmission, from knowledge and justification already possessed. Inference can originate knowledge or justified belief in the sense that the beliefs in question are new to the believer, but apparently not—as the basic sources of knowledge and justification can—from something other than belief, such as perception. Deductive transmission apparently requires

validity; and inductive transmission apparently requires an inductive counterpart of validity, something like a strong relation of support between premises and conclusion. But even when the support is strong, the degree of justification may drop in a way that it need not drop in the deductive case.

As our examples show, to understand the transmission of justification and knowledge we must consider two sorts of conditions: necessary conditions for transmission of knowledge and justification, conditions such that transmission occurs *only if* they are met by an inference; and sufficient conditions, those such that *if* they are met by an inference, then transmission occurs.

It is by and large even harder to specify sufficient conditions than necessary ones. For a sufficient condition must "cover" all the necessary ones: if it does not imply that each of them holds, it leaves out something necessary, and so is not sufficient.<sup>17</sup> Let me simply suggest the sort of thing we must add to what we so far have in order to arrive at sufficient conditions for inferential transmission.

It will help to begin with inductive cases. Might we accept the following *inductive transmission principle*: If, by good inductive reasoning, one infers something from premises which take account of all the relevant evidence, then if one is justified in believing those premises, one is justified in believing the conclusion? Even in the lottery case, in which one holds only one of a million coupons, this condition is plausible for justification. For instance, one may be justified in believing one will lose. Knowledge, however, is different in this respect. For as the lottery example shows, even when the probability is very high, the counterpart of this condition, with knowledge substituted for justification, does not hold.

A different example will show something else about the inductive transmission of justification. Recall the case of a chain of two or more instances of inductive reasoning. Imagine that I enter my house and find evidence of a burglary, such as ransacked drawers. I infer that valuables are missing. From that I infer that the \$20 in my daughter's piggy bank is gone. And from that in turn I infer that my daughter will be upset. At each point I am justified in believing my premise and, apparently make a good inductive inference from it. In most such cases, my justification carries right down the inductive chain from my initial premise to my final conclusion. But it need not. There is a chance that the bank was overlooked and a chance that my daughter will be calm, if only because she is so grateful that important things, such as the teddy bears, are undisturbed. Could it not be that at each step my justification for my conclusion drops in such a way that, unlike my inference that I will lose the lottery, my last inference fails to produce a justified conclusion?

The general point here is that as inductive inference proceeds, the crucially relevant evidence, the evidence one must take into account, may mount up or at least change. For instance, by the time I get to the question of whether my daughter will be upset about the piggy bank, it becomes relevant to note that the teddy bears are unharmed before inferring that she will be upset, whereas this information would not have been relevant if the disappearance of the



piggy bank were the only potential disturbance in the house. We find, then, that the appraisal of inductive chains cannot be accomplished by any simple application of the single inference standard.

But how should we decide what is relevant to drawing a conclusion? And how is one's justification for believing a conclusion affected by ignoring only some of what is relevant? These are hard questions, which I can only partially answer. One positive point is this: whether we are inferentially justified in holding a conclusion we draw depends on many factors, including some *not* expressed in our premises. My believing the premises of an inference may be the origin of my belief and a source of my justification. But there are other relevant factors—such as what I know, or should know, about what will preoccupy the child upon discovering the burglary.

My justification, then, ultimately depends on complex relations among all the relevant factors. We might say that although justification may emerge from a straight inferential line, it will do so only if the line figures in the right kind of *pattern* of related beliefs and available relevant information. Some patterns contain obstacles on the would-be path to justification; others have clear, straight passageways.

By contrast with induction, deduction is more straightforward: if  $p$  entails  $q$  (i.e.,  $q$  is deducible from  $p$ ) and  $q$  entails  $r$ , then  $p$  entails  $r$ ; but if we put 'probabilistically implies' in place of 'entails', we do not get a principle that is invariably true. Suppose the tossing of a weighted (unfair) coin probabilistically implies—i.e., makes more likely than not—the coin's landing heads (toward which the coin is biased); assume that the probability of heads given a toss is just over  $\frac{1}{2}$ . Second, suppose that the coin's landing heads probabilistically implies (also with a probability of just over  $\frac{1}{2}$ ) that (since he will win enough money) Slim will buy a football ticket. It does not follow that his flipping the coin probabilistically implies his buying the ticket. For one thing, the probability of the two independent events, his winning and his buying a football ticket with the money, is the probability of the first *times* that of the second, hence *less* than  $\frac{1}{2}$  (it would be so even if each had a probability as high as 0.70 since  $0.70 \times 0.70$  is only 0.49). If, moreover, we allow for the passage of time, as is inevitable with a person's drawing two or more inferences, then we cannot rule out that intervening factors (such as his detection of the coin's bias as it is tossed) will make it very unlikely that he will accept the money. He may cancel the bet immediately on discovering the unfairness of the coin. These points do not in the least imply that non-entailing grounds for a proposition can never render it highly probable or yield a justified belief of it. The point simply expresses a limitation on inductive principles and chains.

### *Deductive transmission of justification and knowledge*

Let us turn now to deductive transmission. One might think that valid deductive inference is sufficient as well as necessary for transmitting justification



and knowledge. Certainly it commonly does transmit them, for instance when we learn theorems by validly deducing them in doing geometrical proofs.

I am not implying that whenever, in the abstract, there *is* a valid inference, in the sense of 'a set of propositions constituting a valid argument', from something one believes to a conclusion, then one "implicitly" knows the conclusion, or even has situational justification for believing it. If that were so, then simply by knowing the axioms of Euclidean geometry (which, like the parallel axiom—that for any line and any point not on that line, there is exactly one line parallel to the first and passing through the point—are quite simple), one might implicitly know, and be justified in believing, all its theorems, including theorems too complex for one to understand.)<sup>18</sup> For the theorems include all the propositions the axioms entail (an infinite set), and some of those will not even be intelligible to everyone who knows or justifiably believes the axioms. The main issue here is the transmission of justification and knowledge in two ways: first, from justified beliefs, or from beliefs constituting knowledge, to other beliefs *arrived at* by inference; second, from such beliefs to situational justification for propositions that we *could* infer from those we know or are justified in believing.

Even if we restrict attention to transmission of knowledge across inference processes, it is at least not obvious that knowledge is always transmitted across valid deductive inferences.<sup>19</sup> Recall the backfire. Suppose I am sufficiently acquainted with the sound to *know* that it is a backfire. Then, from what I know, it follows that it is not the sound of a firecracker with a similar muffled sound. Imagine that, aware that this follows, I infer that the sound is not that of a firecracker. Do I know this? What if I have no evidence that there is no one around setting off such firecrackers? Perhaps I then do not know this. It may be that from my general experience, the most I am justified in believing is that this alternative explanation of the sound is so improbable that it is irrelevant. But it is still not clear that I know there is no one around setting off such firecrackers.<sup>20</sup>

Thus, it is not clear that we should accept what might be called *the simple deductive transmission principle for knowledge*: that if (at a given time) you validly infer a proposition from an inferential ground you know (and believe it on that ground), then you know this proposition (say that the sound is not that of a firecracker with a similar muffled quality). One might now say that if I do *not* know my inferred conclusion here, this just shows that I did not know in the first place that a vehicle backfired (my premise). But must we say this? It may be equally plausible to say that because one *now* realizes that one's basis for believing this might not have been decisive, one *no longer* knows it, yet did know it in the first place. *If* that is so, it shows something important: that sometimes reflection on our grounds can bring into our purview considerations that weaken our justification for our grounds, or eliminate our knowledge of them, or at least weaken their power to justify our drawing inferences from them.

### *Degrees and kinds of deductive transmission*

Consider a different case. I add a column of fifteen figures, check my results twice, and thereby come to know, and to be justified in believing, that the sum is 10,952. As it happens, I sometimes make mistakes, and my wife (whom I justifiably believe to be a better arithmetician) sometimes corrects me. Suppose that, feeling unusually confident this time, I now infer that if my wife says this is not the sum, she is wrong. From the mathematical truth that the sum is 10,952, it certainly follows that if she says it is not, she is wrong. If it *is* the sum, then if she denies that, she is wrong. But even though I know and am justified in believing that this is the sum, can I, on the basis of my grounds for this belief, *automatically* know or be justified in believing that if she says it is not the sum, she is wrong? (I am assuming that I have no other basis for holding this belief, such as a calculator result that coincides with mine.) That is far from self-evident. To see why, let us focus mainly on the principle as applied to justification—the *simple deductive transmission principle for justification*: If one is justified in believing *p*, then one is also justified in believing any proposition (within one's understanding) that follows from it.<sup>21</sup>

The force of the case is best appreciated if we suppose that my checking just twice is enough to give me only the *minimum* basis for justified belief and knowledge here (the minimum for knowledge seems higher, but that will not affect what follows). Surely I would then not have sufficient grounds for believing that if she says the answer is wrong, she is wrong. Given my background justification for believing that she is the better arithmetician, the justification-threatening prospect this proposition puts before me seems to demand that I have *more* justification than the minimum I do have for my sum, if I am to be justified in believing that if she says the sum is not 10,952, she is wrong.

One way to interpret the example is this. To be justified in believing the proposition that if she says the sum is not 10,952, she is wrong, or to know or justifiably believe this about her, I need grounds for believing this proposition that are good enough not to be outweighed by the supposition that she (the better arithmetician) says that 10,952 is not the sum. In inferring that if she says this is not the sum, she is wrong, I *am* making the supposition that she says it. Of course, I need not *believe* she will say it; but because I am supposing she will (and in a sense envisaging this in considering the proposition that if she says it, she is wrong), I am justified in believing that if she does, she is wrong, *only* if my justification for believing that the sum is 10,952 is good enough to withstand the supposition that she denies it is the sum. My supposing this may also be regarded as implicit in my holding the belief that if she says this, she is wrong, whether I form that belief by inference or not. In either case, under the specified conditions, her justification seems good enough to reduce mine below the threshold of justification, which it just barely reaches.

One might now object that I really do not have justification in the first place for believing that the sum is 10,952. Depending on my arithmetic skills, that might be true if I have checked my sum only once. But suppose that carefully checking three or four times is required to reach the threshold of justification and that I have done this. For *any* reasonable standard of justification or knowledge, there will be a point at which I *just* meet, and do not exceed, that standard, and (again assuming I am justified in believing her to be the better arithmetician) I will then not know or be justified in believing the proposition that if she says the sum is wrong, then she is wrong. (This point concerns situational justification. It is also true that if I infer this further proposition without first getting *additional* grounds for my answer, I would not know it or justifiedly believe it, i.e., have a justified belief of it.)<sup>22</sup>

The example can be varied to make the same point in a different way. If the sum is 10,952, then even if there are two mistakes in my calculations I made, it is *still* 10,952. This may sound strange, but suppose the mistakes cancel each other: one yields a 9 instead of the correct 7, and the other yields a 6 instead of the correct 8. Then an excess of 2 is offset by a shortage of 2.

Now imagine that I again justifiedly believe that the sum is 10,952 and know this. I have been careful enough and have not actually made errors. Still, I have checked only the *minimum* amount necessary for justification. Perhaps simply to test my intuitions about deductive transmission, I might infer that (even) if there are two errors in my calculation leading to the answer, 10,952, the sum is 10,952. Surely I am not justified in believing this and—assuming that the same minimum of checking is sufficient for knowledge—I do not know it (if more checking is required, then the same point will hold for knowledge if we build in the assumption that I just reached the required minimum). My original, minimal justification does not give me situational justification for believing what I infer or adequate grounds for knowledge of that proposition. If I had done *extra* checking, say enough to be adequately justified in believing (or to know) that I made no mistakes, it might be otherwise; but that is not my case.

Still another way to conceive the example is this. One might think of (1) 'If she says the sum is not 10,952, then she is wrong' as equivalent to (2) 'Either she doesn't say this or she says it and is wrong'. Thus, if I am justified in believing (or I know) (2), I am (arguably) justified in believing (or I know) (1). It may seem that I would be justified in believing (2), since such justification can occur in any of these three cases: through my being justified in believing (a) that she will not say this, or (b) that she says it and is wrong, or (c) that at least one of those two things is true. Am I, however, justified in believing any of (a)–(c)? Let us consider them in order.

My justification for believing that the sum is 10,952 is (chiefly) my reasonably careful calculations' indicating that it is. *That* justification does not extend to justifying my believing that (a) my wife will not say that this is false; and it surely does not extend to my believing that (b) she (whom I justifiedly believe to be the better arithmetician) says it is false and is wrong.

Thus, it seems at best unlikely that I should be justified in believing that (c) at least one of these two things is the case.

It is important not to take this cautionary conclusion to be more negative or restrictive than it is. Nothing I have said requires denying that in the imagined case I may also have *some* reason to believe that if she says the sum is not 10,952, she is wrong. But the point is that I would not have *enough* justification for this to know or have a justified belief of it, as I did know and have a justified belief that the sum is 10,952.<sup>23</sup>

Cases of this sort strongly argue for at least two points. First, justification and knowledge *need* not be transmitted through valid inference from known or justifiedly believed premises to belief of a conclusion inferred on the basis of them. Second, situational justification is not automatically transmitted even to propositions clearly entailed by those we are justified in believing—hence the deductive transmission of justification principle is not true.

These negative points should be balanced by another. *Some* degree of (situational) justification—what we might loosely call *some reason for believing*—may automatically transmit: it is not as though I have *nothing* in the way of reason to believe that if she says the sum is not 10,952 then she is wrong; for instance, I did check my calculation with some care. Still, merely having *some* reason to believe does not imply being (unqualifiedly) justified *in* believing, any more than one piece of evidence for a proposition need suffice for knowledge of it.

The sort of failure of transmission I have noted is probably not common for inferences rational persons normally draw, and I stress it because such failure has often been held to be impossible (and is important in dealing with skepticism, as Chapter 13 will show). A *qualified deductive transmission principle for justification* apparently holds: *typically, valid reasoning from justified beliefs transmits justification to its conclusion belief* (this holds for overall justification as well as for some degree of justification, though it allows that there be some degree of diminution in justification across the inference). Similarly, we may apparently affirm a *qualified deductive transmission principle for knowledge*: *typically valid reasoning from known premises transmits knowledge to its conclusion (where the conclusion belief is based on the premise belief(s), as would be normal)*. These principles are of major importance in epistemology. It is difficult to say under just what conditions deductive transmission does not occur, but one can see what some of them are from the points that have emerged here. The general conclusion to draw, however, is that whether one is justified in believing something, or knows it, depends not only on one's specific evidence for it but also on a pattern of factors including one's relation to the proposition itself and one's particular circumstances.

## **Memorial preservation of inferential justification and inferential knowledge**

We should now consider a further point that applies to both deductive and inductive inferential transmission. Imagine that you learn something, say a

theorem, by validly inferring it from something you know, say an axiom. You may remember the axiom *as* your ground; then your memory preserves both your premise and your conclusion. But eventually you may forget your ground, for instance how you proved, and even how *to* prove, a theorem. Similarly, you may forget the testimony or book from which you learned (perhaps by inductive inference partly based on the premise that the book is reliable) that the Battle of Hastings was in 1066. Can you still know and justifiedly hold these now premise-less beliefs?

The answer in both cases is surely that you can. Memory can retain beliefs *as* knowledge, and as justified beliefs, even if it does not retain the original grounds of the relevant beliefs. But because in these instances it does not retain the inferential grounds (or other such as testimonial ones), and no new grounds need be added, it does not necessarily retain the beliefs *as* inferential (or testimony-based). Moreover, when the grounds are not retained and none are added, one might find it at best difficult to indicate *how* one knows, beyond insisting that, say, one is sure one remembers, perhaps adding that one surely did have grounds in the past. So long as one did have adequate grounds and *does* remember the proposition, surely one can know that proposition. (Theories that explain this will be discussed in Chapters 10 and 11). One can also justifiedly believe it, provided one has an appropriate memory of it, say the sense of memorial familiarity that goes with many of the beliefs memory preserves.

This example is another illustration of the point that a belief which is inferential at one time may be non-inferential at another. This may happen repeatedly with the same belief. Long after a belief—for instance, of a theorem—has ceased to be inferential, one could acquire new grounds for it, such as that one has a clear recollection of a mathematical friend's affirming the theorem. One could later forget the new grounds also, and simply remember the theorem or indeed find an altogether new proof of it.

Suppose, however, that one's memory of the theorem is very weak and one has no confidence that one has it right. The result might be that one has merely a belief which not only does not constitute knowledge but also is only weakly justified, if justified at all. It will certainly not be justified if one acquires new evidence that clearly counts strongly against it and nothing happens, such as one's getting new information, to neutralize this hostile evidence. Often, however, the new beliefs, justification, and knowledge we acquire through inference may be retained even when their inferential grounds are long forgotten.

At any given moment in waking life, we have some operative basic source of belief, if only the stream of our own consciousness. As we experience the world around us and our own interactions with it, new beliefs arise, both directly from basic sources and inferentially. As rational beings, we are almost constantly forming beliefs on the basis of other beliefs. We may form these through a process of inference or only through acquiring beliefs

that are only structurally inferential: based on other beliefs and so held for a reason, though not reasoned, since they do not arise from other beliefs (nor are grounded in them) by a process of inference. Both deductive and inductive inferences are common. Both transmit justification and knowledge when they give rise to beliefs on the basis of an inference which meets the appropriate deductive, inductive, and evidential standards.

Among the *transmission principles* that have emerged as plausible are these two broad ones. First, knowledge and justification are inferentially transmitted only if the underlying argument is good. If we start with false or unjustified premises or we unreasonably infer a conclusion from them (i.e., infer it invalidly or in an inductively inadmissible way), it is not to be expected that a belief based on the argument in question constitutes knowledge or is even justified. (This does not, of course, prevent it from having an independent sound basis.) Second, at least typically, if the argument *is* good, (1) situational justification is transmitted and (2) belief justification and knowledge are transmitted provided the subject believes the proposition in question (the conclusion of the inference) on the basis of its premises (the underlying ones).

The kinds of transmission described in the second principle seem to occur quite often, and abnormal conditions such as those described in the column of figures case are apparently not common. Given a normally retentive memory, we have not only a vast store of direct (non-inferential) knowledge and directly justified beliefs, but also a huge variety of indirect knowledge and indirectly justified beliefs. False and unwarranted beliefs arise from some inferences. But from many inferences we learn something new; and in making inferences to propositions that we think best explain something that we take ourselves to know already, we sometimes learn truths that are both new and important. Through inference, then, we often enlarge, strengthen, and develop our body of knowledge and justified beliefs.

## Notes

- 1 Two points may add clarity here. (1) I am talking about beliefs *that* (propositional, not objectual, beliefs). (2) It is perhaps misleading to call propositions objects of beliefs, if only because they can express the content of beliefs—their primary role here—whether or not believing is a relation to a proposition conceived as an object. It could instead be something like a “contentful” property of persons.
- 2 Two points deserve emphasis here. First, I take a belief arrived at by reasoning to be at the time grounded in that reasoning; but it should be noted that the belief can be retained in memory after the premises of the reasoning are forgotten, and in that case we should speak of a belief that is only a formerly reasoned belief. (It is also possible for a belief not arrived at through reasoning to be later grounded in it, as when one at last finds premises to support what one has believed on the basis

of a “hunch.”) Second, one might object to my main distinction here that, from my recognition of the backfire, I *must* have inferred, hence reasoned to the conclusion, that someone drove by. Granted, this recognition is a ground of my belief that someone drove by. Still, I need not do anything that qualifies as drawing a conclusion from the recognition. I did not even stop reading to think about the noise, whereas, in the case of the woodpecker, I focused on the question of whether someone was at the door and, on hearing the distinctive rapid knocking, inferred that it was that of a woodpecker.

- 3 An account of moral perception and its relation to non-moral perception is provided in my ‘Moral Perception and Moral Knowledge’, *Proceedings of the Aristotelian Society Supplementary Volume* 84 (2010), 79–97.
- 4 Granted, the notion of a process of inference is not sharp; sometimes we cannot get enough information about how a belief was formed even to make an educated guess about whether or not it arose from an inference.
- 5 One may wonder what to call an episodically inferential belief just *after* it is formed but when it is held in mind and occurrent. Normally, it would now be in a category to which many beliefs belong: structurally inferential though, unlike most such beliefs, occurrent (it would not *have* to be structurally inferential since it is possible that it ceases to be *based* on the premise(s) from which it inferentially arises).
- 6 The notion of a reason is here understood broadly, so that even a “bad reason” counts. Thus, one could have a belief for a reason even if the reason is a false proposition or one the person is unjustified in believing. Anyone who finds this usage too broad can substitute the notion of a reason or apparent reason. A defense of the broad terminology is given in Chapter 2 of my *Architecture of Reason* (Oxford: Oxford University Press, 2001). If the terminology is too broad, then an episodically inferential belief need not be one held for a reason; for people can draw inferences from propositions that, in the narrow terminology, provide at best an apparent rather than a “real” reason for believing what they believe on the basis of them.
- 7 The distinction between episodically and structurally inferential beliefs and the notion of one belief’s being based on another are discussed in detail in my ‘Belief, Reason, and Inference’, *Philosophical Topics* 14, 1 (1986), 27–65, reprinted in my collection, *The Structure of Justification* (Cambridge: Cambridge University Press, 1993).
- 8 Critics of foundationalism (such as Richard Rorty) have thought that foundationalism is committed to positing intrinsically basic propositions; but we have already seen why even self-evident propositions need not be intrinsically basic, and Chapter 9 will bring out why foundationalism need not posit the latter.



- 9 Bertrand Russell is among a number of philosophers who have at least implicitly denied this: "our knowledge of the physical world is not at first inferential, but that is because we take our percepts [roughly sense-data] to be the physical world . . . adults have got used to the idea that what is really there can only be inferred from what they see." See *An Outline of Philosophy* (London: Allen & Unwin, 1927), chapters 12–13. This is the kind of view criticized by J.L. Austin in *Sense and Sensibilia* (Oxford: Oxford University Press, 1962), esp. chapter 10. Austin subjects A.J. Ayer to criticism on similar counts, especially in the passages reprinted in Huemer, *Epistemology*. Chapter 9 will show that the foundationalism introduced there need not deny the point.
- 10 There is no one way to arrive at a suitable premise, but a typical way would be to count the cats individually up to seventeen, checking to be sure of overlooking none, and to believe on that basis that one has arrived at a total of seventeen cats. There may be ways, however, of passing non-inferentially from counting  $n$  things of a kind  $K$  to the belief that there are  $n$   $K$ s. There are certain premises from which one may infer that if some pets are dogs, then some dogs are pets, e.g. that if, for any  $A$  and  $B$ , if some  $A$ s are  $B$ s, then some  $B$ s are  $A$ s; but no one who can reason thus would find it natural to believe the former on the basis of the latter. On the plausible view we might call *epistemic particularism*, moreover, the former kind of proposition is epistemically prior to the latter, at least in the sense that justification for the former kind is presupposed by justification for the latter kind, but not conversely.
- 11 I say this appears to be so because it is controversial. The issue will be discussed in Chapter 9.
- 12 This is a permissive and formal sense of 'valid' and 'entail', because both apply when the premise set is contradictory or the conclusion is a necessary truth (a truth whose falsity is impossible). For it is impossible that a contradiction be true, hence impossible for a contradictory premise set to be true and the conclusion false; and this is also impossible if the conclusion is a necessary truth and so cannot be false. Usually, we deal with arguments valid in a narrower sense, their premises being both mutually consistent and relevant in subject matter to their conclusion. But nothing said in this book should turn on our using the broad, permissive notion of validity that is standard in formal logic.
- 13 The language of probability seems preferable here to that of defeasibility. For the latter normally implies a liability to loss of justification or knowledge or both, and on that point deductive and inductive reasoning do not differ. If we conjoin to a premise in a valid deductive argument the negation of that premise, the argument and the reasoning expressing



it are still valid (logically indefeasible, one might say); but its premises cannot all be known nor (presumably) justifiedly believed, and in any case they cannot provide good reason to believe the conclusion. They would be further from doing so than in the case of a weak inductive argument. (I thank Claudio de Almeida for questioning the text in a way that brought out this kind of problem.)

- 14 For an informative discussion of knowledge ascriptions highly pertinent to whether knowledge admits of degrees and supportive of the view suggested here, see chapter 2 of Jason Stanley, *Knowledge and Practical Interest* (Oxford: Clarendon Press, 2005).
- 15 Note that the claim that one would not know here is most plausible when 'know' is stressed; this is one reason for the plausibility of contextualism about claims to know. For extensive discussion of lottery cases, see John Hawthorne, *Knowledge and Lotteries* (Oxford: Oxford University Press, 2004). Contextualism will be discussed in Chapter 10.
- 16 The point here is associated with what is called the *lottery paradox*, introduced into the literature by Henry E. Kyburg, Jr. and widely discussed. See his *Epistemology and Inference* (Minneapolis: University of Minnesota Press, 1983).
- 17 Since a sufficient condition implies all of the conditions that are minimally necessary, i.e., are the (possibly complex) conditions individually necessary and jointly sufficient for the phenomenon to occur, some have wondered how a sufficient condition can fail to be a necessary one as well. The answer is that it can imply something more that is not necessary. For instance, taking a letter to the postbox by car, although sufficient for getting it there, is not necessary for this, since it implies something not necessary for getting it there—driving.
- 18 I am assuming that being justified in believing *p* entails being able to understand it. This is not to deny that one could have justified beliefs, say of a set of axioms, that *are* a justification for a proposition, *p*, which one does not understand, so that someone who had the same justified axiomatic beliefs and understood *p* *would* be justified in believing *p*. The distinction, then, is between possessing a set of beliefs, or being in a state, that *is* a justification for *p*, and being justified in believing *p* in the usual sense. Some who speak of propositional justification may not observe this distinction. I thank Scott Hagaman for calling my attention to a usage that does not observe the distinction.
- 19 I mean, of course, the non-trivial kind of valid inference, having consistent premises none of which is equivalent to the conclusion. From inconsistent premises anything may be validly derived. If, for example, we start with a premise consisting of (1) some proposition, *p*, and its negation, not-*p* (i.e., start with a contradiction), we may infer that (2)

either  $p$  or  $q$ , for any proposition  $q$  we like (on the ground that if  $p$  holds, then either it or anything whatever holds). But we may now bring in (3) not- $p$  (since we have it as well as  $p$  as one conjunctive element in our premise); and not- $p$ , together with (2), entails  $q$ . Our arbitrarily chosen proposition,  $q$ , is thus validly derived.

- 20 On some views, a central feature of knowledge is that the belief in question is justified in a way that allows one to rule out, or itself in some way rules out, relevant alternatives. Valuable discussion of this issue is provided by Alvin I. Goldman, 'Discrimination and Perceptual Knowledge', *Journal of Philosophy* 73 (1976), 771–91. The issue is addressed, sometimes indirectly, in Chapter 10.
- 21 This principle should be qualified if the idea is to be maximally plausible. It is not obvious what the best formulation might be, but we should at least specify that the person in question can understand the entailment from  $p$  to the proposition for which the person has justification by transmission. There is no reason to think that because I am justified in believing  $p$  I should be justified in believing  $q$ , which it entails, when I cannot for the life of me understand how the former entails the latter. I am also for the most part simplifying by speaking only of transmission of situational justification as opposed to belief justification. For transmission of that, as of knowledge, it is presumably required that the person believe the entailed proposition on the basis of the entailing one.
- 22 This column of figures example has generated considerable discussion in the literature. For detailed critical discussion of my case see Catherine Canary and Douglas Odegard, 'Deductive Justification', *Dialogue* 28 (1989), 305–20; Richard Feldman, 'In Defense of Closure', *Philosophical Quarterly* 45 (1995), 487–94; and Peter D. Klein, 'Skepticism and Closure: Why the Evil Demon Argument Fails', *Philosophical Topics* 23, 1 (1995), 213–36. For my replies published to date see 'Justification, Deductive Closure, and Skepticism', *Dialogue* 30 (1991), 77–84; and 'Deductive Closure, Defeasibility, and Skepticism: A Reply to Feldman', *Philosophical Quarterly* 45 (1995), 494–9. (This paper discussed the example construed as appealing to a subjunctive such as 'If she were to say the sum is not  $n$ , she would be wrong'.) For related treatments of the transmission problem see Fred Dretske, 'Epistemic Operators', *Journal of Philosophy* 67 (1970), 1007–23; Robert Nozick, *Philosophical Explanations* (Cambridge, MA: Harvard University Press, 1981), chapter 3; Gilbert Harman, *Change in View* (Cambridge, MA: MIT Press, 1986), esp. chapters 1–4; and Marian David and Ted A. Warfield, 'Knowledge-Closure and Skepticism', in Quentin Smith (ed.), *Epistemology: New Essays* (Oxford: Oxford University Press, 2008).

- 23 No precise notion of justified belief fits all the contexts in which we speak of it; but when we speak of justification unqualifiedly we usually have in mind a kind and degree such that it is reasonable, overall, for the person to hold the belief. This is perhaps a kind and degree such that, if the proposition in question is true (and there are no special problems of the kind to be considered in Chapter 10), then the belief constitutes knowledge. Even this high degree of justification, of course, is still best conceived as *prima facie* rather than indefeasible or in any sense absolute.

## 9 The architecture of knowledge

- **Inferential chains and the structure of belief**

- Infinite inferential chains*

- Circular inferential chains*

- **The epistemic regress problem**

- Infinite epistemic chains*

- Circular epistemic chains*

- Epistemic chains terminating in belief not constituting knowledge*

- Epistemic chains terminating in knowledge*

- **The epistemic regress argument**

- **Foundationalism and coherentism**

- **Holistic coherentism**

- Patterns of justification*

- A coherentist response to the regress argument*

- **The nature of coherence**

- Coherence and explanation*

- Coherence as an internal relation among cognitions*

- Coherence, reason, and experience*

- Coherence and the a priori*

- Coherence and the mutually explanatory*

- Epistemological versus conceptual coherentism*

- Coherence, incoherence, and defeasibility*

- Positive and negative epistemic dependence*

- **Coherence and second-order justification**

- The process versus the property of justification*

- Beliefs, dispositions to believe, and grounds of belief*

- Justification, knowledge, and artificially created coherence*

- **Moderate foundationalism**

- The role of coherence in moderate foundationalism*

- Moderate foundationalism and the charge of dogmatism*

## 9 The architecture of knowledge

On the mountain in the distance before me, I see the huge oak and tulip trees swaying, with their leaves turned upward revealing the lighter green of their undersides. Waves of green, light and then dark, seem to cross the surface of the upper region from west to east as the leaves show the colors of their inner and outer sides. Parts of the hillside seem almost to breathe in and out as the trees bend away from me and back. It is a familiar sight, and I immediately realize that there is a wind. My belief that there is a wind is based on my belief that the trees are swaying. It is also justified on the basis of that belief. And if I know that there is a wind, I know it on the basis of my belief that they are swaying. In each case, one belief is inferentially based on another.

To what extent does this relation in which one belief is based on another represent the structure of our belief systems as a whole? The question is especially pertinent to epistemology as applied to the common cases in which our beliefs constitute knowledge. Might perceptual beliefs, for instance, be a foundation on which many others are inferentially built? Or are perceptual beliefs just a stopping place on the way to something yet more basic, or perhaps merely a place where we usually stop pursuing further premises, though we might go on seeking them and find deeper grounds that support perceptual beliefs?

These questions represent perennial issues, and we shall see many versions of the *foundationalist view*—the classical position on them—and various opposing theories. The questions also lead us, as often happens in epistemology, into questions about the nature of mind as well as questions directly about justification and knowledge. This is to be expected when the central topic is the structure of knowledge and justification; for knowledge is apparently constituted by belief, and, in epistemology, justification is important chiefly in connection with belief. It is appropriate, then, to begin an exploration of the structure of knowledge and justification with some major points about the structure of a person's body of beliefs.

### Inferential chains and the structure of belief

As in discussing inference, it is useful to call the kind of inferential belief, justification, and knowledge just illustrated *indirect*. For one has such beliefs,

justification, and knowledge only on the basis of, and thereby *through*, other beliefs, justification, or knowledge. By contrast, my belief that the trees are swaying is direct. I believe this simply because I see it, not on the basis of something else I believe.

### *Infinite inferential chains*

The natural picture just sketched can be challenged. Perhaps all our beliefs could be *indirect*. If they could be, could not all justification of belief, and all our knowledge, be indirect? An adequate epistemology requires answers to these questions about the structure of a body of belief, justification, or knowledge. In exploring them, I will talk above all about knowledge and justification, and especially about knowledge. But what we know (propositionally) we believe; and the kind of justification epistemology is chiefly concerned with is that of belief. The structure of my knowledge and justification, then, is chiefly that of a certain body of my beliefs.

I am not talking about knowledge in the abstract, as we sometimes do. We speak, for instance about the extent of “human knowledge.” Some of this knowledge is solely in books, and not remembered by anyone. Thus, some scientific knowledge might be constituted by propositions no one actually believes: they are *available* to us if we need them, but not objects of actual belief. They might be accessible to us by our simply consulting, or drawing straightforward inferences from, scientific literature. We *can* talk about the structure of such knowledge in the abstract, say about whether all the propositions of scientific knowledge can be *systematized* by certain basic laws of physics and chemistry treated as explaining or entailing (or both) the others. Then these basic laws would be something like geometrical axioms, and the other laws, like theorems, would be derivable from the basic laws. But that is not our topic here (though it will be discussed in some detail in Chapter 12). We are here exploring how people’s beliefs may actually be structured.

Consider a simple example. When I am being very cautious, my belief that the trees are swaying could be based on my belief that I have a visual impression of swaying. Could the latter belief also be based on another one? What might that be? Might I now believe that it *seems* to me that I have a visual impression of swaying, and base my belief that I have that impression on this new belief? This is doubtful. I cannot base one belief on another simply because I want to.

This example shows that the view that what we believe, and certain relations between our beliefs, are entirely under the direct control of our wills—a strong version of *doxastic voluntarism* (voluntarism about belief)—is a mistake. Suppose, for instance, that I want to believe someone’s testimony. If it seems false, I cannot make myself believe it just by willing myself to believe it. Similarly, I generally cannot, at will, change *relations* among beliefs I hold. I lack, for instance, direct voluntary control over what my beliefs are based on; if I already know first-hand, from my sense of weakness, that I am gravely ill, I cannot, simply by willing it, *base* my belief of this on someone’s

testimony that it is so. This holds even if I believe the testimony and take it to be confirmatory.<sup>1</sup>

Even if one cannot base one belief on another at will, it might still seem that a sequence of beliefs, each based on the next, could go on without limit. But could I, for instance, believe what seems the next proposition in the evidential series, the involuted proposition that it appears to me that it seems to me that I have a visual impression of swaying? I suppose I could (though not simply at will). Still, I do not see that I would now come to hold anything on the *basis* of believing this strange proposition.

Suppose, however, that I did come to hold, on the basis of this involuted proposition, that it seems to me that I have an impression of swaying. I cannot in this way manufacture an *inferential chain* of beliefs—a chain in which each belief is based on the next—running to infinity. Nor do I already have an infinite set of appropriate beliefs as raw material waiting to be brought to consciousness—if indeed I *can* have an *infinite* number of beliefs (particularly outside mathematics, in which it may seem that I can have an infinite number corresponding to the series 2 is even, 4 is even, 6 is even, etc.).<sup>2</sup>

### *Circular inferential chains*

So far, however, I have ignored another way in which it might be thought to be possible that every belief is indirect: by virtue of lying not at the origin of an infinite chain, but instead in a circular chain. Imagine that I could hold one belief on the basis of a second and a second on the basis of a third, and so on, until we come full circle and get to a belief I hold on the basis of the first. Then all my beliefs would be indirect, yet I need not have infinitely many. To assess this, recall my belief that there is a swaying. Might there be a circular chain of beliefs here? For instance, could my belief that it appears to me that it seems to me that I have a visual impression of swaying be based on my belief that there is a swaying? This is far from clearly possible.

Suppose for the sake of argument that I do have a circular chain of beliefs, each based on the next. This raises a problem. First, there is good reason to think that (a) one belief is based on a second only if the second is at least in part causally responsible for (one's holding) the first. For instance, if I believe there is a wind, on the basis of my believing that the trees are swaying, then I believe that there is a wind, at least in part *because* I believe that the trees are swaying. Second, there is good reason to think that (b) if one thing is in part causally responsible for a second and the second is in part causally responsible for a third, then the first is in part causally responsible for the third. But together these two points imply that (c) in a circular chain of beliefs, each based on the next, every belief is in part causally responsible for, and thus a partial cause of, itself. That seems impossible. To see why, let us explore how such a circle might go in a simple case.

Imagine a circle of three beliefs, each based on the next. (1) I believe there is a wind. I believe this on the basis of (2) my believing there is a swaying of

the trees; I believe that there is this swaying, on the basis of (3) my believing I have an impression of such swaying; and, coming full circle, I believe that I have this impression, on the basis of believing there is a wind. This case would be a *circular causal chain*, one whose last link is connected to its first in the same way that each is connected to its successor. For, given point (a), belief (1) is in part causally responsible for belief (3), and, given point (b), (3) is in part causally responsible for (1). This implies, however, given (b), that (1) is in part causally responsible for itself. That is apparently impossible. The belief would be holding itself up by its bootstraps.

If we accept the lesson of this *bootstraps problem* (as perhaps some will not) we must conclude that circular causal chains of this kind are not possible and so there cannot be a circular chain of beliefs each based on the next. For on the highly plausible assumptions, (a) through (c), this would have to be a circular causal chain. (We have not assumed that the imagined chain implies that some belief must be *based* on itself, only that such chains imply a belief's being in part causally responsible for itself; the basis relation implies more than a causal connection.)

It might seem that a wheel is a model of a circular causal chain of the relevant kind and that something must therefore be wrong with the reasoning just noted. Consider a wheel standing on the ground in a line running east and west, and imagine the wheel having eight equal sections and an axle, each section consisting of a pie-slice segment with its apex at the axle. Does each section not support the next, so that each rests on the others and ultimately (in the eighth link) on itself?

If we distinguish between the relation of *being connected with* and that of *supporting*, the answer no longer seems clear. Granted that if one section is connected to a second, it *will* support the second (to some degree) *if* a force is applied to the second in the direction of the first. But a wheel with eight such connected sections can exist in empty space with no such forces acting on it. Mere connectedness between segments does not imply any actual support relations, only a readiness to enter them.

Consider, then, the realistic case in which the wheel is on the ground. Gravity exerts a downward force on the entire wheel. Here, however, the ground supports the entire wheel, and each segment of the wheel that has a segment above it supports that segment, with the two top sections (whose common seam, we may assume, runs straight up from the center of the wheel to its highest point) being the only ones plausibly said to support each other directly. But notice that each of the top sections supports the other *with respect* to a different force. There is a westward force in the case of the western section's support of the eastern one (which would fall backwards to the east if disconnected from its western counterpart because all the seams become unfastened); and there is an eastward force in the case of the eastern section's support of the western one (which would fall backwards to the west if disconnected from its eastern counterpart because all the seams become unfastened).



Each top section of the wheel, then, pulls on the other in the opposite direction, with the result being a balance. In no case do we get a force in *one* direction that goes fully around the circle with the result that any section supports itself in that same direction. The forces on the two top sections are, as described in physics, equal and opposite.

Returning to the case of belief, there the support in question—the kind of cognitive support given by one belief to a second that is based on it—is also in one direction. It is, in good part, support with respect to three dimensions: conviction, explanation, and memory. Consider this cognitive force in relation to a common case, that of a conclusion belief being based on a belief of a premise, such as a point made by a respected friend in favor of the conclusion. My premise belief tends to increase or buttress my conviction in my conclusion belief, to explain (in part and in some broadly causal way) why I hold that belief, and to help me remember my conclusion. Even if one takes circular causation to be possible, one should grant that this is not the kind of support relation that a belief may be plausibly thought to bear *to itself*.

Self-support on the part of beliefs might also seem possible if one conceives self-evident propositions as “self-justified” and takes their self-justification to be a kind of self-support. One might think, for instance, that a belief of a self-evident proposition can be in part causally responsible for itself and in that way support itself. But ‘self-justified’ is misleading: a self-evident proposition is not justified *simpliciter*; propositions are justified or not *for* someone, and self-evident ones are justified for those who adequately understand them and on the basis of that understanding. To call a belief of a self-evident proposition self-justified seems at best an inaccurate way of saying that such a proposition is not normally believed because one believes *something else*. It is of course also not believed because one believes *it*. Normally, one believes it because one grasps the appropriate conceptual relation(s) it expresses. In any case, our concern is beliefs in general, not just beliefs of self-evident propositions.

On balance, then, it is reasonable to conclude not only that we *have* direct beliefs, such as beliefs ascribing colors to things before us and beliefs of self-evident propositions, but also that we could not have *only* indirect beliefs. Apparently, neither infinite nor circular chains of indirect beliefs are possible for us.

## The epistemic regress problem

Is knowledge like belief in this, so that some of it is direct, or could *all* our knowledge be indirect, that is, based on other knowledge we have? It may seem that this is possible, and that there can be an infinite *epistemic regress*—roughly, an infinite series of knowings each based on the next.

It is especially likely to appear that indirect knowledge need not always be based on direct knowledge, if one stresses that, very commonly, ‘How do you know?’ can be repeatedly answered, and one then supposes that we

stop answering only for practical reasons having to do with our patience or ingenuity. Let us explore this issue by assuming for the sake of argument that there is indirect knowledge and seeing what this implies.

Assume that a belief constituting indirect knowledge is based on knowledge of something else, or at least on a further belief. The further knowledge or belief might be based on knowledge of, or belief about, something still further, and so on. Call this sequence an *epistemic chain*. It is simply a chain of beliefs with at least the first constituting knowledge, and each belief linked to the previous one by being based on it.

It is often held that there are just four possible kinds of epistemic chain. Two kinds are unanchored and do not end; two kinds are anchored and do end. First, an epistemic chain might be infinite, hence entirely unanchored. Second, it might be circular, hence also unanchored. Third, it might end with a belief that is not knowledge, and thus (figuratively speaking) be anchored in sand. Fourth, it might end with a belief that constitutes direct knowledge, and thus be anchored in bedrock. Our task is to assess these chains as possible sources of knowledge or justification. This is a version of the epistemic regress problem.

### *Infinite epistemic chains*

The first possibility is difficult to appreciate. Even if I could have an infinite number of beliefs, how would I ever know anything if knowledge required an infinite epistemic chain? To know, and thus to *learn*, the simplest kind of thing, such as that there is a green field before me, I would apparently have to know an infinite number of things.

It is doubtful that, given our psychological make-up, we *can* know, or even believe, infinitely many things. It might seem that we can have an infinite set of arithmetical beliefs, say that 2 is larger than 1, that 3 is larger than 2, and so forth. But surely for a finite mind there will be some point or other at which the relevant proposition cannot be grasped (the limiting point might be different for different people or even the same person at different times). Imagine the “largest” proposition a supercomputer could formulate after years of work. It could easily be too complex to understand or so cumbersome that one could not even take in a formulation of it. One would be unable to remember enough about the first part of it when one gets to the end; one could thus never understand the whole thing. What we cannot understand we cannot believe; and what we cannot believe we cannot know.<sup>3</sup>

Even if we could have infinite sets of beliefs, however, infinite epistemic chains apparently could not account for all, and probably not for any, of our knowledge. In the case of some beliefs, such as the belief that if some dogs are pets, some pets are dogs, I cannot even find *any* belief I hold that yields another link (a belief this one seems to be based on). The proposition is luminously self-evident, and it is difficult even to imagine a further proposition I would consider a good premise on the basis of which I would believe it if I

thought I needed a premise for it. Thus, I find it unclear how this belief could be grounded, as knowledge, by any epistemic chain, much less by an infinite one.

In any event, how might infinite epistemic chains help us account for any other knowledge (or justified belief)? Notice that many kinds of infinite chain are possible. No one has provided a plausible account of what kind might generate justification or knowledge. But some restrictions are badly needed. For any proposition, an infinite chain can be imagined (in outline) that may be claimed to provide support for the proposition. Thus, even for a proposition one believes to be obviously false, one would find it easy to imagine beliefs to back it up; and though one could not continue doing this to infinity, one could nonetheless claim that one *has* the infinite set required to support the original belief.

Take the obviously false proposition that I weigh at least 500 pounds. I could back up a belief of this by claiming that if I weigh at least 500.1 pounds, then I weigh at least 500 (which is self-evident), and that I do weigh at least 500.1 pounds. I could “defend” *this* by appeal to the propositions that I weigh at least 500.2 pounds, and that if I do, then I weigh at least 500.1. And so forth, until the challenger is exhausted. A chain like this can be infinite; hence, no matter how ridiculous a proposition I claim to know, there is no way to catch me with a claim I cannot back up in the same way. Given such resources, anything goes. But nothing is accomplished.

### *Circular epistemic chains*

The possibility of a circular epistemic chain as a basis of knowledge has been taken more seriously. It might seem that if there cannot be a circular causal chain of indirect beliefs, each based on the next, then there cannot be a circular epistemic chain either. But perhaps knowledge can be based on premises in a way that differs from the way belief is based on them; perhaps, for instance, my *knowledge* that there is a wind could be somehow based on my belief that the leaves are swaying, even though my *belief* that there is a wind is not based on any further belief. We would then have a circle of knowledge, but not of belief, and no causal bootstraps problem. If this is possible, it may turn out to be important. But how realistic is it?

Does any of our knowledge really emerge from circular epistemic chains? Let us try to go full circle. I know there is a wind. I know this on the basis of the swaying of the trees. Now I think I know they are swaying because I *see* them sway. But it might be argued that my seeing this is only the causal basis of my *belief* that they are swaying, and I just do not notice that it is only on the basis of, say, my knowledge that I have a visual impression of swaying that I *know* they are swaying. Perhaps. But how far can this go?

I do not see how to go full circle, unless I *think up* propositions I do not originally believe, hence do not originally know. If I do not originally believe

them, then I (originally) have no *justified* belief or knowledge of the premise they constitute, and thus no belief appropriate to serve as a link in the epistemic chain or play any supporting role toward my original knowledge.

Suppose, however, that I do think up a suitable set of evidential propositions, come to know them, and make my way full circle. Suppose, for instance, that I get as far as knowledge that it seems to me that I have a visual impression of swaying. Might I know this on the basis of knowing that there is a wind (the first link)? How would knowledge that there is a wind justify my belief that it seems to me that I have a visual impression of tree swaying? I apparently know introspectively, not perceptually or inferentially, that I have the impression of swaying. Other difficulties also beset the circular approach. But these problems alone cast sufficient doubt on it to suggest that we consider the remaining options.

### *Epistemic chains terminating in belief not constituting knowledge*

The third possibility for the structure of epistemic chains, that an epistemic chain terminates in a belief which is not knowledge, can be best understood if we recall that in discussing the transmission of knowledge, we noted both source conditions and transmission conditions. If the third possibility can be realized, then knowledge can originate on the basis of a belief of a premise that is not known. On the basis of believing that there is a swaying, for example, I might know that there is a wind, even though I do not know that there is a swaying. The regress is thus stopped by grounding knowledge on something else, but not in the *way* it is normally grounded in experience or reason.

Is this possible? In one kind of case it is not. Suppose that (in foggy conditions) I simply guess that what I see is a swaying of trees, but happen to be right. Might I then know there is a wind anyway, provided there is? Surely not; knowledge cannot be grounded in such guesswork, even when the guess is correct.

Imagine, however, that although I do not know there is a swaying, I do hear some sounds that might indicate swaying, and I make an *educated* guess and am thereby justified, to some significant degree, in believing that there is. If, on the basis of this somewhat justified belief that there is a swaying, I now believe that there is a wind, and there is, do I know this?

The answer is not clear. But that would be no help to proponents of the third possibility, who claim that knowledge can arise from belief which does not constitute knowledge. For it is equally unclear, and for the same sort of reason, whether my guess that there is a swaying is *sufficiently* educated—say, in terms of how good my evidence is—to give me (a weak kind of) knowledge that there is a swaying. If it is clear that my guess is *not* sufficiently educated to yield this knowledge, then I also do not know there is a wind. If it is clear

that the guess is educated enough, I apparently do know that there is a wind, but my knowledge would be based on other knowledge, hence would not realize the third possibility.

Notice something else. In the only cases in which the third kind of chain is at all likely to ground knowledge, there is a degree—perhaps a substantial degree—of justification. If there can be an epistemic chain which ends with belief that is not knowledge only because the chain ends, in this way, with justification, then it appears that we are at least in the general vicinity of knowledge. We are at most a few degrees of justification away. The sand has turned out to be rather firm; it is at least close to being firm enough to support knowledge.<sup>4</sup>

### *Epistemic chains terminating in knowledge*

The fourth possibility is the one apparently favored by common sense: epistemic chains end in direct knowledge—in the sense that they have direct knowledge as their last link. That knowledge, in turn, is apparently grounded (anchored, if you like) in experience or in reason, and this non-inferential grounding explains how it is (epistemically) direct: it arises, directly, from perception, memory, introspection, or reason (or indeed from testimony, provided this has an appropriate ultimate grounding in at least one of the first four).

The ground-level knowledge just described could not be inferential; otherwise the chain would not end without a further link. To illustrate, normally I know that there is a swaying just by seeing that there is. Hence, the chain grounding my knowledge that there is a wind is anchored in my perception.

Such experientially or rationally grounded epistemic chains may differ in many ways. Here are four. They differ in *composition*, in the sorts of beliefs constituting them. They differ in the kind of *transmission* they exhibit; it may be deductive, inductive, or combine both deductive and inductive links. Epistemic chains also differ in their ultimate grounds, the *anchors* of the chains, which may be experiential or rational; and epistemic chains may vary in justificational *strength*, the degree of justification they give to the initial belief.

Different proponents of the fourth possibility have held various views about the character of the *foundational knowledge*, that is, of the beliefs constituting the knowledge that makes up the final link of the epistemic chain that is anchored in experience or reason. Some philosophers, for instance, have thought that the appropriate beliefs must be infallible, or at least indefeasibly justified. But this is not implied by anything said here. All that the fourth possibility requires is *direct* knowledge, knowledge not based on other knowledge (or on justified belief).

Direct knowledge need not be of self-evident propositions, or constituted by indefeasibly justified belief. Introspective beliefs illustrate this. The proposition that I am now thinking about knowledge is not self-evident. It is not

even *self-evident* to *me*. First, it is evident to me, not in itself, as is the proposition that if some dogs are pets then some pets are dogs, but on the basis of my conscious experience. Second, since I realize that my reflections can sometimes merge into daydreaming, I do not even consider it rock-solidly true in the way I do self-evident propositions. But surely I do have direct knowledge of the proposition.

## The epistemic regress argument

What we have just seen suggests a version of the *epistemic regress argument*. It starts with the assumption that:

- 1 If one has any knowledge, it occurs in an epistemic chain.

Epistemic chains are understood to include the special case of a single link, such as a perceptual belief, which constitutes knowledge by virtue of being anchored directly (non-inferentially) in one's experience.<sup>5</sup>

The argument then states that:

- 2 The only possible kinds of epistemic chain are the four mutually exclusive kinds just discussed: the infinite, the circular, those terminating in beliefs that are not knowledge, and those terminating in direct knowledge.

Its third, also restrictive, premise is that:

- 3 Knowledge can occur only in the fourth kind of chain.

And the argument concludes that:

- 4 If one has any knowledge, one has some direct knowledge.<sup>6</sup>

A similar argument was advanced by Aristotle, and versions of this regress argument have been defended ever since.<sup>7</sup>

As proponents of the argument normally understand (1), it implies that any given instance of indirect knowledge depends on at least one epistemic chain for its status *as* knowledge. So understood, the argument clearly implies the further conclusion that any indirect knowledge a person has *epistemically depends on*, and cannot be knowledge apart from, an appropriate inferential connection, via some epistemic chain, to some direct knowledge that the person has.

Given this dependence assumption, the regress argument would show not only that if there is *indirect* knowledge, there is direct knowledge, but also that if there is indirect knowledge, that very knowledge is *traceable* to some direct knowledge as a foundation for it. One could trace an item of indirect knowledge to some premise for it, and, if there is a premise for that,

to the next premise, and so on, until the chain is anchored in a basic source of knowledge.

A similar argument applies to justification. We simply speak of *justificatory chains* and proceed in a parallel way, substituting justification for knowledge; and we arrive at the conclusion that if one has any justified beliefs, one has some directly justified beliefs. Similarly, if one has any indirectly justified belief, it exhibits *justificational dependence* on an epistemic chain appropriately linking it to some directly justified belief one has, that is, to a foundational belief.

## Foundationalism and coherentism

These two sets of conclusions constitute the heart of the position called *epistemological foundationalism*. The first set, concerning knowledge, may be interpreted as the thesis that the structure of a body of knowledge (such as yours or mine) is foundational, in which this is taken to imply that any indirect (hence non-foundational) knowledge there is depends on direct (and thus in a sense foundational) knowledge. The superstructure, one might say, rests on the foundations. The second set of conclusions, regarding justification, may be interpreted as the thesis that the structure of a body of justified beliefs is foundational, where this is taken to imply that any indirectly (hence non-foundationally) justified beliefs there are depend on directly (thus in a sense foundationally) justified beliefs.

In both cases, different foundationalist theories may diverge in the kind and degree of dependence they assert. A strong foundationalist theory of justification might hold that indirectly justified beliefs derive *all* their justification from foundational beliefs; a moderate theory might maintain only that the former would not be justified apart from the latter, and the theory might grant that other factors, such as coherence of a belief with others one holds that are *not* in the chain, can add to its justification.

None of the foundationalist theses I have stated says anything about the *content* of a body of knowledge or of justified belief, though proponents of foundationalism usually specify, as René Descartes does in his *Meditations on First Philosophy* (first published in 1641), what sorts of content they think appropriate. Foundationalism, as such, thus leaves open what, in particular, is believed by a given person who has knowledge or justified belief and what *sorts* of propositions are suitable material for the foundational beliefs. I want to talk mainly about foundationalism regarding knowledge, but much of what I say can be readily applied to justified belief.

Foundationalism has been criticized on a number of points. Let us focus on the most important objections that stem from the most prominent alternative theory of the structure of knowledge, *coherentism*. There are many versions of coherentism, including some that seem to be based mainly on the idea that if an epistemic circle is large enough and sufficiently rich in content, it can generate justification and account for knowledge. But we have seen

serious difficulties besetting circular chains. I therefore want to formulate a more plausible version of coherentism.

The central idea underlying coherentism is that the justification (justifiedness) of a belief depends on its coherence with other beliefs one holds. The *unit of coherence*—roughly, the range of the beliefs that must cohere in order for a belief among them to derive justification from their coherence—may be as large as one's entire set of beliefs (though of course some may figure more significantly in producing the coherence than others, say because of differing degrees of closeness to one another in their subject matter).

The variability of the unit of coherence would be accepted by a proponent of the circular view, but the thesis I want to explore differs from that view in not being *linear*: it does not construe justification or knowledge as emerging from an inferential line going from premises to that conclusion, and from other premises to the first set of premises, and so on, until we return to the original proposition as a premise.

On the circular coherentist view, no matter how wide the circle, there is a *line* from any one belief in a circular epistemic chain to any other. In practice we may never *trace* the entire line, as by inferring one thing we know from a second, the second from a third, and so on, until we re-infer the first. Still, on this view, for every belief that constitutes knowledge, there is such a line, however long it may be. Thus, the kinds of problems we encountered earlier regarding circular epistemic chains must be resolved (as I doubt they can be) if the view is to be sustained.

## Holistic coherentism

Coherentism need not be linear. It may be holistic. To see how a holistic theory of knowledge (and justification) works, consider a question that evokes a justification. John wonders how I know, as I sit reading, that the wind is blowing. I say that the leaves are rustling. He then asks how I know that Sally is not just making this noise by walking in the high grass. I reply that the high grass is too far away. He now wonders whether I can distinguish rustling leaves from the sound of a quiet car on the pebbled driveway. I reply that what I hear is too much like a whisper to be the crunchy sound of pebbles under tires.

### *Patterns of justification*

In giving this kind of justification, I apparently go only one step along the inferential line: just to my belief that the leaves are rustling. For my belief that there is a wind is based on this belief about the leaves. After that, I do not even mention anything that this belief, in turn, is based on. Rather, I defend my beliefs as appropriate in terms of an entire pattern of mutually cohering beliefs I hold. And I may cite many different parts of the pattern. For instance, I might have said that walking through high grass sounds different



from windblown leaves. On the coherentist view, then, beliefs representing knowledge do not have to lie in a grounded chain; they fit a coherent pattern, and their justification emerges from their fitting that pattern in an appropriate way. If this justificational *fittingness* seems like a kind of foundation for justification, that is because in a structural sense it is; but coherentists hold that it differs from classical foundations in being a property of the belief system rather than something, such as perception, that is outside it.

Consider a different sort of example. A gift is delivered to you with its card apparently missing. The only people you can think of who send you gifts at this time of year live in Washington and virtually never leave, but this is from Omaha. That origin does not cohere well with your hypothesis that it was sent by your Washington benefactors, the Smiths. Then you open it and discover that it is frozen steak. You realize that this can be ordered from anywhere. But it is not the sort of gift you would expect from the Smiths. A moment later you recall that you recently sent them cheese. You suppose that they are probably sending something in response. Suddenly you remember that they once asked if you had ever tried frozen gourmet steaks, and when you said you had not, they replied that they would have to give you some one of these days.

You now have a quite coherent pattern of beliefs and might be justified in believing that it was they who sent the package. If you come to believe this on the basis of the pattern, you presumably have a justified belief. When you at last find their card at the bottom of the box, then (normally) you would *know* that they sent the package.

The crucial things to notice here are how, initially, a kind of *incoherence* with your standing beliefs prevents your justifiedly believing your first hypothesis (that the box came from the Smiths) and how, as relevant pieces of the pattern developed, you became justified in believing, and (presumably) came to know, that the Smiths sent it. Arriving at a justified belief, on this view, is more like answering a question in the light of a whole battery of relevant information than like deducing a theorem by successive inferential steps from a set of luminous axioms.

### *A coherentist response to the regress argument*

It is important to see how, using examples like those just given, holistic coherentism can respond to the regress argument. It need *not* embrace the possibility of an epistemic circle (though its proponents need not reject that either). Instead, it can deny the premise that there are only the four kinds of possible epistemic chains so far specified. There might be a fifth: a chain terminating with belief that is *psychologically direct*, yet *epistemically indirect* (or, if we are talking of coherentism about justification, *justificationally indirect*). This in effect grants foundationalists that they are right about human psychology, while insisting that they are wrong about epistemology. Let me explain.

The idea is that although a terminal, direct belief is not *psychologically* based on any other, as when it is inferentially grounded on another, its justification nonetheless *is* based on other beliefs. Hence, the last link is, as belief, direct, yet, as knowledge, indirect, not in the usual sense that it is inferential but in the broad sense that the belief constitutes knowledge only by virtue of receiving support from other knowledge or belief. This belief is *psychologically foundational but epistemically dependent*. Its justification depends on a pattern of supporting beliefs.

To illustrate all this, consider again my belief that the trees are swaying. It is psychologically direct because it is grounded, causally, in my vision and is not inferentially based on any other belief. Yet (the coherentist might argue) my *knowledge* that there is such a movement is not epistemically direct. It is epistemically, but not inferentially, based on the coherence of my belief that there is a rustling with my other beliefs, presumably including many that represent knowledge themselves. It is thus knowledge *through*, but not by inference from, other knowledge—or even through justified beliefs. The knowledge is therefore epistemically indirect. Hence, it is at best misleading to call the *knowledge*, as opposed to the belief expressing it, direct at all.

This coherentist view grants, then, that the belief element *in* my knowledge is non-inferentially grounded in perception and is in that sense direct. But this is just a kind of psychological directness: there is no belief through which I hold the one in question in the way that I hold a conclusion belief on the basis of premise beliefs. But there are beliefs through which the belief constitutes knowledge: those with which it coheres even though it is not based on them. The basis relation between beliefs and the counterpart premise–conclusion relation between propositions are simply not the only producers of coherence.

One could insist that if a non-inferential, thus psychologically direct, belief constitutes knowledge, this *must* be direct knowledge. But the coherentist would reply that in that case there will be two kinds of direct knowledge: the kind the foundationalist posits, which derives from grounding in a basic experiential or rational source, say perception or reflection, and the kind the coherentist posits, which derives from coherence with other beliefs and not from being based on those sources. Why not classify the directness of knowledge in terms of what it *evidentially* depends on and the directness of belief in terms of what it *psychologically* depends on? This is surely a plausible response.

Is the holistic coherentist trying to have it both ways? Not necessarily. Holistic coherentism can grant that a variant of the regress argument holds for belief, as the only kind of belief chain that it is psychologically realistic to attribute to us is the kind terminating in direct (non-inferential) belief: there are no infinite or circular belief chains in any finite mind. But even on the assumption that knowledge is constituted by (certain kinds of) beliefs, it does not follow that direct belief which is knowledge is also direct *knowledge*.

Thus, the coherentist is granting *psychological foundationalism*, which

says (in part) that if we have any beliefs at all, we have some direct ones, yet denying epistemological foundationalism, which says that, assuming there is any knowledge at all, there is knowledge which is epistemically (and normally also psychologically) direct. Holistic coherentism may grant experience and reason the status of psychological foundations of our entire structure of beliefs. But it gives them no place, independently of coherence, in generating justification or knowledge.<sup>8</sup>

The idea that a belief might be psychologically direct (so non-inferential) but epistemically indirect is not the only basis on which coherentists and others can challenge the idea that knowledge occurs only in epistemic chains terminating in knowledge. It might also be held (even apart from coherentism) that knowledge might occur in a chain that is indirect in *both* ways and even terminates in a belief that is not knowledge. Imagine a sister and brother, Carla and Jan, who believe Santa Claus brings Christmas presents. Aware of their having been threatened with a bad report to Santa Claus, Carla asks Jan whether there will be presents this Christmas. He truly replies, "Yes, Mama said Santa will bring them." If there will be presents, we might suppose that he knows this. But his premise is false given that the presents will be provided by his parents, not Santa. He appears, then, to have knowledge based on a false belief—though a useful falsehood—one in which indeed the epistemic chain terminates.<sup>9</sup>

Here a defender of the regress argument might appeal to the distinction made above between the *basis of a belief that constitutes knowledge* and the *basis of the knowledge it constitutes*. The coherentist may hold that the latter basis is coherence with other beliefs even if the knowledge is constituted by a non-inferential belief; by contrast, the defender of the regress argument may hold that even if the *belief* constituting knowledge is inferential the knowledge it constitutes need not be inferential. We would thus have direct knowledge that *p* as the terminal link in an epistemic chain, even though the belief that *p* is inferential.

How might that be? One plausible hypothesis is that although the boy's belief that there will be presents is *inferential*, being based on his believing Mama said that Santa will bring them, his *knowledge* that there will be is *non-inferential*. He knows it because his belief that there will be is appropriately grounded in a fact that guarantees its truth: it is produced (in part) by her intention to give the presents, which in turn produces her testimony that Santa would bring presents, which in turn produces his belief that there will be presents. The crucial point is that his belief that there will be presents derives from a fact that guarantees there will be and does so in such a way as to make the mistaken premise—mistaken only as to who will bring the presents—epistemically harmless.

In Chapters 10 and 11 more will be said about the conception of knowledge that makes this foundationalist response plausible, but here I am mainly concerned with the epistemological resources of foundationalism and coherentism. Supposing the regress argument is sound, either view can deal with

useful falsehood in question by using the distinction between the basis of a belief that constitutes knowledge and the basis of the knowledge it constitutes. The hypothesis suggested here favors a kind of foundationalism, but it should not be assumed that coherentists cannot provide an alternative hypothesis to account for the children's knowledge.

## The nature of coherence

As I have described holistic coherentism, it avoids some of the major problems for linear coherentism. But there remain serious difficulties for it. First, what is coherence? Second, what reason is there to think that coherence *alone* counts toward the justification of a belief, or toward its truth, as it must in some way if it is to give us the basis of a good account of knowledge?

It turns out to be very difficult to explain what coherence is. It is not mere mutual consistency, though inconsistency is the clearest case of *incoherence*. Two propositions having nothing to do with each other, say that  $7 + 5 = 12$  and that carrots are nourishing, are mutually consistent but do not exhibit coherence.

## Coherence and explanation

Coherence is sometimes connected with *explanation*. Certainly, if the Smiths' sending the package explains why the card bears their names, then my belief of the first proposition coheres with my belief of the second (other things being equal). What explains something makes it understandable; and *making understandable* is a coherence-generating relation between propositions (as well as between other kinds of things).

*Probability* is also relevant to coherence. If the probability of the proposition that the Smiths sent the steaks is raised in the light of the proposition that I sent them cheese, this at least counts in favor of my belief of the first cohering with my belief of the second. But how are we to understand the notions of explanation and of probability? Let us consider these questions in turn.

Does one proposition (genuinely) explain another so long as, if the first is (or at least is assumed to be) true, then it is clear why the second is true? Apparently not; for if that were so, then the proposition that a benevolent genie delivered the box explains why it arrived. In any event, if that proposition did explain why the box arrived, would I be justified in believing it because my believing it coheres with my believing that I know not what other source the box might have come from? Surely not.

Even if we can say what notion of explanation is relevant to understanding coherence, it will remain very difficult to specify when an explanatory relation generates *enough* coherence to create justification. For one thing, consider cases in which a proposition, say that Jill hurt Jack's feelings, would, if true, very adequately explain something we believe, such as that Jack is

upset. Believing Jill did this might cohere well with his being upset, but that would not, by itself, justify our believing it. There might be too many possible competing explanations we might just as well accept. There might also be far better ones, such as that his computer has been stolen.

Similar points hold for probability. Not just any proposition I believe which, if true, would raise the probability of my hypothesis that the gift is from the Smiths will strengthen my justification for believing that it is. Consider, for example, the proposition that the Smiths send such gifts to all their friends. Suppose I have no justification for believing this, say because I have accepted it only on the basis of testimony which I should see to be unreliable. Then, although the proposition, if true, raises the probability of my hypothesis (since I am among their friends) and (let us assume) coheres with what I already believe, I am not thereby entitled to believe it, and my believing it will not add to my justification for believing that the Smiths sent the box.

It might be replied that this belief about the Smiths' habits does not cohere well with *all* that I believe, such as that people do not generally behave like that. But suppose I believed nothing about the Smiths' or anyone's habits of gift-giving that conflicts with the Smiths' being so generous, and—unjustifiably—I believed the Smiths to be both generous and rich. Then there might be a significant degree of coherence between my belief that the Smiths send gifts to all their friends and my other beliefs; yet my forming the belief that they give gifts to all their friends still would not strengthen my justification for my hypothesis that the steak is from them.

### *Coherence as an internal relation among cognitions*

These examples suggest the second problem. So far as we do understand coherence, what reason is there to think that by itself it generates any justification or counts toward truth? Whatever coherence among beliefs is, it is an *internal* relation: whether it holds among beliefs is a matter of how those beliefs (including their propositional content, which is intrinsic to them) are related *to one another*. It is not a matter of anything outside one's system of beliefs, such as one's perceptual experience. Now why could there not be numerous equally coherent systems of beliefs that are mutually incompatible, so that no two of them can be without falsehood? If there can be, why should my having one of these coherent systems provide any reason to think *my* beliefs, rather than those of someone with one of the "opposing" systems, are justified or represent knowledge?

This is part of what might be called the *isolation problem*: the problem of explaining why coherent systems of beliefs are not readily isolated from truth, and thus do not contain knowledge, which implies truth. There is also a problem of explaining why there is not a similar isolation from justification, which seems in some way to point toward truth, roughly in the sense that what justifies a belief "indicates" its truth, and indicates it in proportion to

the degree of justification. Why should coherence by itself imply that any of the cohering beliefs is justified or constitutes knowledge, when both justification and knowledge point toward truth as something external to the belief system? It is not as though coherentists could count on the implication's being guaranteed by God; and nothing else seems to assure us of it.

Consider a schizophrenic who thinks he is Napoleon. If he has a fully consistent story with enough interlocking details, his belief system may be superbly coherent. He may even be able to explain quite coherently why there are coherent belief systems that conflict with his, such as those of his psychiatrists. If coherence alone generates justification, however, we must say that each system is equally well justified—assuming their belief systems are exactly as coherent as his. We need not attribute knowledge to any of the systems; any of them might contain only falsehoods.

But is it plausible to say that a system of beliefs is highly justified even when there is no limit to the number of radically different yet equally justified belief systems—even on the part of other people with experience of many of the same things the beliefs are about—that are incompatible with it in this thoroughgoing way? The question is especially striking when we realize that two equally coherent systems, even in the same person at different times, could differ not just on one point but on *every* point: each belief in one system might be opposed by an incompatible belief in the other.<sup>10</sup>

To see the significance of the possibility of multiple coherent mutually incompatible systems of belief, recall the plausible assumption that a well-justified belief may reasonably be considered *true*. If, however, the degree of justification of a belief is entirely a matter of its support by considerations of coherence, no degree of justification by itself can carry any greater presumption of truth than is created by the same degree of support from coherence on the part of a belief of the contradictory proposition. Thus, if “Napoleon” (unlike his historical namesake) has a sufficiently coherent set of beliefs yielding justification of his belief that he won the Battle of Waterloo, this belief may be as well-justified as his psychiatrists’ belief that he was not even born at the time.

If this coherentist picture of justification is correct, is there any reason to think that a belief supported solely by considerations of coherence is true or even justified? And if “Napoleon’s” and the psychiatrists’ belief systems are equally coherent, how can we justify our apparently quite reasonable tendency to regard their belief systems as more likely to represent truths, and on that count more likely to contain knowledge, than his?

Granted, the psychiatrists’ belief that he was born long after the battle coheres with our beliefs. But why should our own beliefs be privileged over equally coherent conflicting sets? And why should agreement even with nearly everyone’s beliefs, say about Napoleon’s being dead, be a factor, unless we assume that some element other than coherence, such as perception or memory, confers justification without drawing on coherence? If coherence is the only source of justification, it is not clear how perception or memory

or introspection contributes to justification. Even what seems the highest degree of justification, such as we have for beliefs of luminous self-evident truths and for simple introspective beliefs, provides us no presumption of truth or knowledge.

### *Coherence, reason, and experience*

This brings us to a third major problem for coherentism: how can it explain the role of experience and reason as apparent sources of justification and knowledge? Certainly experience and reason *seem* to be basic sources of them. Coherentists themselves commonly *use* beliefs arising from these sources to illustrate coherent bodies of beliefs that are good candidates for knowledge. How can holistic coherentism explain the relevance of these sources to justification or knowledge?

Why is it, for instance, that when I have a vivid experience of the kind characteristic of seeing a green field, I am apparently justified (though *prima facie*, rather than indefeasibly, justified), simply by that experience, in believing that there is a green field before me? And why do I seem so very strongly justified, simply on the basis of my rational grasp of the proposition that if some dogs are pets then some pets are dogs, in believing this?

One thing a coherentist might say here is that in fact many of our beliefs are *causally* and non-inferentially based on perception or reason; and given these similarities of origin, it is unsurprising that they often cohere with one another. Hence, although we do not, and need not, infer propositions like those just cited from any others that might provide evidence for them, they *do* cohere with many other things we believe, and this coherence is what justifies them.

### *Coherence and the a priori*

This response by way of associating the coherence of beliefs with their causal basis is more plausible for perceptual beliefs than for beliefs of simple self-evident a priori truths, at least if coherence is construed as more than consistency and as related to explanation, probability, and justification. For notice that, unlike the proposition that there is a green field before me, the proposition that if some dogs are pets then some pets are dogs can be justified for me even if it does not explain, render probable, or justify anything else I believe. Nor is it obvious that anything else I believe need explain, render probable, or justify my believing this proposition. Why is coherence required for my justification? I may *have* other beliefs that cohere with this one, but my justification for it does not seem to *derive* from such coherence. Yet my belief of this proposition is justified to about as high a degree as is any belief I have.

By contrast, the proposition that there is a green field before me perhaps does cohere, in a way that might serve coherentism, with other things I



believe: that there is grass there, that I am on my front porch, and so on; and there seem to be some explanatory and probability relations among these propositions. For instance, that there is a green field before me adds to the probability that I am on my porch; and that I am on that porch partly explains why I see a green field.

A coherentist might respond to the difference just indicated by qualifying the coherence view, applying it only to beliefs of empirical, rather than a priori, propositions.<sup>11</sup> This move could be defended on the assumption that propositions known a priori are necessarily true and hence are not appropriately said to be made *probable* by other propositions, or to be explained by them in the same way empirical propositions are explained. In support of this it might be argued that although we can explain the *basis* of a necessary truth and thereby show *that* it holds, still, since it cannot fail to hold, there is no explaining *why* it, *as opposed to something else*, holds.

This is plausible but inconclusive reasoning. We may just as reasonably say that we can sometimes explain why a necessary truth holds and in doing so explain why a contrasting proposition is false. Imagine that someone mistakenly takes a certain false proposition to be a theorem of logic and cannot see why a closely similar, true proposition is a theorem. If we now prove the correct one step by step, with accompanying examples, we might thereby explain why this theorem, as opposed to the other proposition, is true.

So far as explanation is central to coherence, then, coherentism apparently owes us an account of knowledge of at least some necessary truths. But suppose that it can account for knowledge of *some* necessary truths. There remain others, such as simple, luminously self-evident ones, for which it cannot offer anything plausibly said to explain why they hold, or any other way of accounting for knowledge of them as grounded in coherence.

Consider how one might explain why, if it is true that Jane Austen wrote *Persuasion*, then it is not false that she did. If someone did not see this, it would probably not help to point out that no proposition is both true and false. For if one needs to have the truth of such a clear and simple instance of this general truth explained, one presumably cannot understand the general truth either. But suppose this is not so, and that one's grasp of the general truth is somehow the basis of one's seeing the particular truth that instantiates it. Then the same point would apply to the general truth: there would apparently be nothing plausibly said to explain to one why *it* is true.

### *Coherence and the mutually explanatory*

It might now be objected that the general truth that no proposition is both true and false, and the instances of it, are *mutually explanatory*: its truth explains why they hold, and their truth explains why it holds; and this is the chief basis of their mutual coherence. But is it really possible for one proposition to explain another *and* the other to explain it? If what explains why the grass is wet is that there is dew on it, then the same proposition—that there is



dew on it—is not explained by the proposition that the grass is wet (instead, condensation explains why it is dewy).

Reflection on other purported examples of mutual explanation also suggests that two propositions cannot explain each other. It might seem that a man could say something because his wife did, and that she could say it because he did. But notice how this has to go to make good sense. One of them would have to say it first to cause the other to. But then we would have a case in which something like this occurs: her saying it explains why he says it, later (this could be so even if her saying it is explained by her believing he *thinks* it). His saying it earlier than she does might still explain her saying it. But then the fact that he says it *at a given time* does not both explain and get explained by her saying it at some particular time.

When we carefully specify what explains something, we seem to find that the latter, carefully specified, does not explain the former. In the case in which she says something because he did, earlier, and he says it because she did, earlier than he did, we would have a kind of *reciprocal* explanation, wherein a *kind* of thing, here spousal affirmation, explains and is explained by another thing of the same kind. But this is not a *mutual* explanation, wherein the very *same* thing explains and is explained by a second thing.<sup>12</sup>

Perhaps mutual explanation of the kind the coherentist apparently needs—as opposed to reciprocal explanation and other sorts involving two-way relations—is somehow possible. But until a good argument for it is given, we should conclude that even if an explanatory relation between propositions is sufficient for a belief of one of the propositions to cohere with a belief of the other, coherentism does not provide a good account of knowledge of self-evident truths.

If coherentism applies only to empirical beliefs, however, and not to beliefs of a priori propositions, then it is not a *general* theory of justification or knowledge and leaves us in need of a non-coherentist account of a priori justification (and knowledge). In any case, it would be premature to conclude that coherentism does account for empirical justification. Let us return to the perceptual case.

### *Epistemological versus conceptual coherentism*

It might seem that we could decisively refute the coherence theory of justification by noting that one might have only a single belief, and that this lone belief might still be justified. For there would be a justified belief that coheres with no other beliefs one has. But could one have just a single belief? Could I, for instance, believe that there is a green field before me, yet not believe, say, that it has any vegetation? It is not clear that I could; and foundationalism does not assume this possibility, though the theory may easily be wrongly criticized for implying it.

Foundationalism is in fact consistent with *one* kind of coherentism, namely, a coherence theory of the acquisition, function, and nature of

concepts—for short, the *coherence theory of concepts*, on which *concepts are what they are partly in relation to one another*, and a person acquires concepts, say of (physical) objects and shapes, and of music and sounds, only in relation to one another and must acquire an entire set of related concepts in order to acquire any concept. The concept of an object in some way includes that of shape (if only the notion of something bounded), as that of music includes the concept of sound. This may be why any object must have some shape or other, and why anything that makes music produces some sound. One cannot (fully) acquire object concepts without acquiring some shape concepts, or (fully) acquire the concept of music without acquiring that of sound.

If the coherence theory of concepts is sound, foundationalists must explain how it squares with their epistemology. The central point they may appeal to is a distinction between (justificatory) *grounding conditions* for belief and *possession conditions* for it. What grounds a belief in such a way as to justify it or render it an item of knowledge is largely independent of *what* other beliefs one must have, and what concepts one must have, to be able to *hold* the first belief. Perhaps I cannot believe that music is playing if I do not have a concept of sound; I may even have to believe sounds with a certain structure to be occurring. And perhaps I could not have acquired these and other relevant concepts one at a time. Indeed, it may be (as suggested in Chapter 7) that at least normally we cannot acquire concepts without acquiring some knowledge or justified belief. Still, what it *is* that justifies a belief can be a matter of how the belief is grounded; it need not be a matter of the coherence conditions required for *having* the belief.

If, however, coherence relations are essential for holding a belief at all, they are on that ground necessary for, and—in ways that will soon be apparent—important in understanding, the belief's being justified. The point here is simply that we cannot treat conditions for having a belief at all as doing the more specific job of grounding its justification. By and large beliefs can be possessed without being justified, and there is usually a good distance between meeting the conditions for simply having a belief and meeting the standards for justification in holding it.<sup>13</sup>

### *Coherence, incoherence, and defeasibility*

We must directly ask, then, whether my justification for believing that there is a green field out there (when I experience its presence) *derives* from the coherence of the belief with others. Let us first grant an important point by focusing on a line of reasoning that may lead many philosophers to think it does derive from coherence. Suppose this visual belief turns out to be *incoherent* with a second belief, such as that one is standing where one seems to see the field around one yet *feels* no grass on the smooth ground beneath one and can walk right across the area without feeling any. Then the first belief may *cease* to be justified. Incoherence, then, defeats my justification.

This defeating role of incoherence is important, but it shows only that our justification is *defeasible*—liable to being outweighed (overridden) or undermined—should sufficiently serious incoherence arise. It does not show that justification is produced by coherence in the first place, any more than a wooden cabin's being destroyed by fire shows that it was produced by the absence of fire. Where I feel no grass beneath my feet, then, the justification of my visual belief is outweighed: my better justified beliefs, including the conviction that a field must have a certain texture, make it more reasonable for me to believe that there is *not* a field here.

A major lesson that emerges here is that we cannot tell what the basis of something is just from the range of things that outweigh it, much less conclude that this basis *is* the absence of the things that destroy it. Incoherence is absent when there are mutually irrelevant beliefs as well as when there are mutually coherent ones. Mutual irrelevance between two sets of beliefs certainly does not make one of them a justificational or epistemic basis for the other.

Two important questions arise here. First, could incoherence outweigh justification of a belief in the first place if we were not *independently* justified in believing something to the effect that a proposition incoherent with certain other ones is, or probably is, false? Second, are the other relevant propositions not precisely the kind for which, directly or inferentially, we have some degree of justification through the basic experiential and rational sources? Foundationalists are likely to answer the first negatively and the second affirmatively.

There is also a different kind of defeat of justification: our justification can be simply *undermined*. We can cease to be justified in believing a proposition, though we do not become justified in believing it false, as one does when counter-evidence demands a contrary belief. Suppose I cease to see a bird on a branch when, without obscuring my line of sight to the bird, I move six feet to my left. This could justify my believing that I might be hallucinating. This belief is incoherent with, and thereby undermines the justification of, my visual belief that the bird is there, though it does not by itself justify my believing that there is *no* bird there.

Again, I am apparently justified, independently of coherence, in believing that my seeing the bird there is incoherent with my merely hallucinating it there. It seems that coherence has the role it does in justification largely because *some* beliefs are justified independently of it.

### *Positive and negative epistemic dependence*

Examples like these show that it is essential to distinguish *negative epistemic dependence*—which is a form of defeasibility—from *positive epistemic dependence*—the kind beliefs bear to the sources from which they *derive* any justification they have or, if they represent knowledge, derive their status as knowledge. The defeasibility of a belief's justification by incoherence

does not imply, as coherentists have commonly thought, that this justification positively depends on coherence. If my well is my source of water, I (positively) depend on it. The possibility that people could drain it does not make their non-malevolence part of my source of water, or imply a (positive) dependence on them, such as I have on the rainfall. Moreover, it is the rainfall that *explains* both my having the water and its level. Its not being drained does not explain this.

So it is with perceptual experience as a source of justification. Foundationalists need not claim that justification does not depend negatively on anything else, for as we have seen they need not claim that justification must be indefeasible. Its vulnerability to defeat can be construed as a kind of dependence. A belief's justification is, then, not *completely* independent of the justification of other beliefs, actual or hypothetical. But negative dependence does not imply positive dependence. Justification can be defeasible by incoherence, and thus outweighed or undermined should incoherence arise, without owing its existence to coherence in the first place.

## Coherence and second-order justification

There is something further that may be considered supportive of coherentism, and assessing it will clarify both coherentism and justification. If we set out to *show* that a belief is justified, we have to cite propositions that cohere with the one in question, say the proposition that there is a green field before me. In some cases, these are not even propositions one already believes. Often, in defending a belief, one forms new beliefs, such as the belief one acquires, in moving one's head, that one can vividly see the changes in perspective that go with seeing a bird on a branch.

### *The process versus the property of justification*

More importantly, these new, back-up beliefs are especially appropriate to the *process of justifying* one's belief; and the result of that process is (a kind of) *showing* that the original belief is justified, together (in typical cases) with one's forming a certain *second-order belief*—so called because it is a belief about a belief (such as a perceptual one) which is not itself about any other belief. In this case the second-order belief is to the effect that the first-order belief is justified. Thus, coherence is important in showing that a belief is justified and is in *that* sense an element in a typical kind of process of justification.

The moment we reflect on this point, however, we may wonder why the beliefs appropriate to *showing* that a belief is justified are required for its *being* justified in the first place. There is no good reason to think they are. Indeed, why should our simply having a justified belief imply even that we are (situationally) justified in holding beliefs appropriate to showing that it is justified? It would seem that just as we can be virtuous even if we do not

know how to defend our good character against attack or even show that we have good character at all, we can have a justified belief even if, in response to someone who doubts that we do, we could not show that we do.

Justifying a second-order belief is a sophisticated process. The process is particularly sophisticated if the belief concerns a special property such as the justification of the original belief. Simply being justified in a belief about the color of an object is a much simpler matter.

Confusion is easy here because of how we often speak of justification. Consider the question of how a simple perceptual belief “is justified.” The very phrase is ambiguous. For all it tells us, the question could be ‘By what process, say of reasoning, might the belief be justified?’ or, by contrast, ‘In virtue of what is the belief justified (possessed of the property of justifiedness)?’ These are very different questions. But much talk about justification makes it easy to conflate them. A belief said to be “justified” could be one that *has* justification or one that *has been* justified; and asking for someone’s justification could be either a request for justifying factors or an invitation to recount the process by which the person has in fact justified the belief.

Does coherentism have any plausible argument, not grounded in the mistakes just noted, for (positive) dependence of perceptual justification on coherence? I do not see that it does, though given how hard it is to discern precisely what coherence is, we cannot be confident that no direct argument is forthcoming. One could, for instance, point to the oddity of saying things like, ‘I am justified in believing that there is a green field there, but I cannot justify the belief’. Coherentists might think this is odd because they tend to hold that if one has a justified belief, one can justify it by appeal to other beliefs one holds that cohere with it. But look closely. Granted that, commonly, in asserting something I suggest that I *can* justify it in some way or other (particularly if the belief I express is not grounded in a basic source), here it could well be my *asserting* that my belief is justified, rather than its being so, that gives the appearance that I must be able to give a justification for the belief if it is justified. In asserting that I am justified, after all, I have not, or not merely, *expressed* a first-order belief, something a normal child of three can do; I have *ascribed* first-order justification to my belief. That requires some sophistication. More important, even foundationalists who hold that we are typically directly justified in, say, perceptual beliefs may deny that normally we are directly justified in these sophisticated ascriptions of justification. To hold that there are non-inferentially justified beliefs does not in the least commit one to holding that ascriptions of justification itself are thus justified.

### *Beliefs, dispositions to believe, and grounds of belief*

To be sure, when I say that there is a green field before me, I (as an educated adult) can give a justification: that I see it. But first, *giving* a justification is

not equivalent to claiming that one has it. The first *cites* a justifier and need not employ the concept of justification; the latter employs that sophisticated concept and need not cite a justifier. Second, note that *before* the question of justification arises I need not even *believe* that I see the field. That question leads me to focus on my circumstances, in which I first had a belief solely about the *field*, not about my own perceptual relation to it.

To be sure, when I said there is a green field before me, I did have a *disposition*, based on my visual experience, to form the belief that I *see* the field, and this is largely why, in the course of justifying that belief, I then form the further belief that I *do* see it. But a *disposition to believe* something does not imply one's already having a *dispositional belief* of it: here I tend to form the belief that I see the field if, as I view it, the question whether I see it arises; yet I need not have subliminally believed this already.

Thus, the justification I offer for my belief that there is a green field before me is not by appeal to coherence with other beliefs I *already* hold—such as that I saw the field and heard the swishing grass beneath my feet—but by reference to a basic source, sensory experience. It is thus precisely the kind of justification that foundationalists are likely to consider appropriate for a non-inferential belief. Indeed, one consideration favoring foundationalism about both justification and knowledge, at least as an account of our justificational practices in everyday life (including much scientific practice), is that typically we cease offering justification or defending a knowledge claim precisely when we reach one or more of the basic sources.<sup>14</sup>

Suppose, however, that I would be dumbfounded if asked, in clear daylight, what justifies me in believing there is a green field before me. Would it follow that I am not justified? No, for I might be simply unable to marshal my quite ample justificatory resources. Coherentism offers no good argument to show that being justified requires being able to show that one is, any more than having good character entails being able to show that one has it.

### *Justification, knowledge, and artificially created coherence*

There is one further point here. If coherentism regards justification as deriving from coherence alone, then it accords *no* justificatory weight to experiential or rational grounding except insofar as they contribute to coherence. Our examples cast much doubt on this view.

Consider a related implication of coherentism. If I seek the best justified body of beliefs possible—surely a rational goal—then I am free to consider adopting, or to manipulate my brain to cause my forming, an entirely new system of beliefs. Would its coherence alone guarantee that it contains justified beliefs? It might contain none of the experiential and a priori beliefs I now have; and for all coherence requires it may entirely lack beliefs based on experience or reason.

A superbly coherent system of beliefs I might acquire could even run counter to my experience. Even if I *see* a square field of green grass before

me, I might coherently believe that there is an oval field of brown shrubbery there, since my other beliefs might support this. I could, for instance, coherently believe that when I seem to see green grass I am having a hallucination caused by brown shrubbery. There is no limit to the number of beliefs for which I might be able thus to rationalize away the states and events that it is natural to call the *evidence of the senses*.

We are apparently incapable of changing our belief systems in this way. But suppose that we could do so by properly setting a neurological machine to instill an optimally coherent set of beliefs and remove the rest. Would that be rational from the point of view of maximizing the justification of one's beliefs? I doubt this, particularly if, in seeking justification, we aim, as we normally do, at discovering or retaining *truths*.

A coherentist might reply that if we are talking not only about justification but also about *knowledge*, then we must give some special role to beliefs (and perhaps dispositions to believe) grounded in experience and reason, for if we ignore these sources we cannot expect our justified beliefs to be true, hence cannot expect them to constitute knowledge.<sup>15</sup> Now, however, we face an artificial separation between what justifies a belief and what is plausibly taken to count toward its truth. If, because it implies truth, knowledge must in some way reflect experience or reason, should not justification, which also seems to count toward truth, also reflect them? May we plausibly suppose that what justifies a belief may in no way count towards its truth?

It is not reasonable to separate justification and knowledge in this way (even though in some ways they are very different); nor have coherentists generally thought that it is (though some have held a justification-based coherence theory of truth of a kind to be discussed in Chapter 11). Often, what motivates asking for a justification of a belief is doubt that it is true; and if so, then the view that what justifies a belief has no tendency whatever to count toward its truth seems plainly mistaken. Moreover, if we can know *a priori*, as I believe may be possible (and will explore in Chapter 13), that perceptual and rational grounding of beliefs count, in some way, toward their truth, why may we not know equally well that they count toward justifying beliefs?

## Moderate foundationalism

There is far more to say about both foundationalism and coherentism. But if what has emerged here is on the right track, then the problems confronting coherentism are more serious than those confronting foundationalism. The most serious problems for foundationalism are widely taken to be, first, the difficulties of specifying source conditions for justification and knowledge and, second, of accounting, on the basis of those sources and plausible transmission principles, for all that we seem to know. The first of these problems is addressed in Part One, which describes the basic sources and illustrates how they generate direct—though not indefeasible—knowledge, and direct



(though again not generally indefeasible) justification. The second problem is treated in Chapter 8, which indicates how, even without actual inferences, knowledge and justification can be transmitted from beliefs which are justified, or represent knowledge, by virtue of being grounded in the basic sources, to other beliefs. Both problems are difficult, and neither has been completely solved here. But enough has been said to clarify along what lines they can be dealt with.

### *The role of coherence in moderate foundationalism*

Still another problem for foundationalism is the difficulty of accounting for the place of coherence in justification. But this is not a crippling difficulty for the kind of foundationalism I have described, which need not restrict the role of coherence any more than is required by the regress argument. Indeed, although (pure) coherentism grants nothing to foundationalism beyond perhaps its underlying psychological picture of how our belief systems are structured, foundationalism can account for some of the insights of coherentism, for instance the point that we need a coherence theory of the acquisition and function of concepts.

More positively, foundationalism can acknowledge a significant role for coherence in relation to justification and can thereby answer one traditional coherentist objection. I have in mind a kind of *moderate foundationalism*: a foundationalist view of knowledge or justification which (1) takes the justification of foundational beliefs to be at least typically defeasible; (2) is not *deductivist*, that is, does not demand that principles governing the inferential transmission of knowledge or justification be deductive (i.e., require entailment as opposed to probability as a condition for transmission); and (3) allows a significant role for coherence by requiring, not that inferentially justified beliefs derive *all* their justification from foundational ones, but only that they derive enough of it from the latter to remain justified if any other justification they have were eliminated.<sup>16</sup> Some versions are more moderate than others, but the most plausible ones give coherence at least two roles.

The first role moderate foundationalism may give to coherence, or strictly speaking to incoherence, is negative: incoherence may defeat justification or knowledge, even of a directly justified (foundational) belief, as when my justification for believing I may be hallucinating prevents me from knowing, or even remaining justified in believing, that the green field is before me. (If this is not ultimately a role for coherence itself, it is a role crucial for explaining points stressed by coherentism.)

Second, moderate foundationalism can employ a principle commonly emphasized by coherentists, though foundationalists need not grant that the justification or truth of the principle is based on coherence and will tend to treat it as a transmission principle accounting for generation of inferential justification or as a combinatorial principle applying to the simultaneous testimony of sources of non-inferential justification. I refer to an *independence*



*principle*: that the larger the number of independent mutually consistent factors one believes (with some justification) to support (or to constitute evidence for) the truth of a proposition, the better one's justification for believing it (other things being equal). This principle can explain, for instance, why my justification for believing that the box of steaks is from the Smiths increases as I acquire new beliefs, each of which I believe independently supports that conclusion.<sup>17</sup> In part, the idea is that evidential relations generate coherence; hence by giving the former a justificatory role, foundationalism can account for a good many of the cases in which coherence appears to yield justification. That appearance may be due not to coherence itself, but to its basis in the kinds of elements foundationalism takes to ground justification.

Similar principles consistent with foundationalism can accommodate other cases in which coherence enhances justification, say those in which a proposition's explaining, and thereby cohering with, something one justifiably believes tends to confer some degree of justification on that proposition.

### *Moderate foundationalism and the charge of dogmatism*

Moderate foundationalism contrasts with *strong foundationalism*, which, in one form, is deductivist, takes foundational beliefs as indefeasibly justified, and allows coherence at most a minimal role. To meet these conditions, strong foundationalists may reduce the basic sources of justification to reason and consciousness. The easiest way to do this is to take the skeptical view (considered in Chapter 13) that our only justified beliefs are either a priori or introspective.

Moreover, since strong foundationalists are committed to the indefeasibility of foundational justification, they would not grant that incoherence can defeat the justification of foundational beliefs. They would also refuse to concede to coherentism, and hence to any independence principle they recognize, any more than a minimal positive role, say by insisting that if a belief is supported by two or more independent cohering sources, its justification is increased at most incrementally, that is, at most by bringing together the justification transmitted separately from each relevant basic source.<sup>18</sup>

By contrast, what moderate foundationalism denies regarding coherence is only that it is a basic source of justification: coherence by itself is not sufficient for justification. Thus, the independence principle does not apply to beliefs or other sources that have *no* justification. At most, it allows coherence to raise the level of justification originally drawn from other sources to a level *higher* than it would have if those sources were not mutually coherent.

Similarly, if inference (at least valid or inductively strong inference) is a basic source of coherence (as some coherentists seem to have believed), it is not a basic source of justification. It may lead to justification, as when one strengthens one's justification for believing someone's testimony by inferring the same point from someone else's. But inference *alone* does not generate justification: I might infer any number of propositions from several

I already believe merely through wishful thinking; yet even if I thus arrive at a highly coherent set of beliefs, I have not thereby increased my justification for believing any of them. My premises, based in the way they are on desire, are ill-grounded.

At this point it might occur to one that the main problems faced by coherentism could be solved by taking *coherence with experience* to be required by coherentism as a condition for the coherence of a body of beliefs of the kind we normally have. This is, to be sure, *not* how coherence is characteristically understood by coherentists; they typically take it to be a relation among beliefs or their propositional contents or other items that may be said to be true or false, or some combination of these.<sup>19</sup> Might it be, however, that leading coherentists misrepresent the resources of their own theory? Could they claim, for instance, that if my visual experience contains an appearance of a printed page, then my believing there is one before me coheres with my experience and is thus justified?

If we think this, we must ask how a coherentist view that gives a crucial epistemological role to coherence of beliefs with experience differs from a moderate foundationalism. One would, after all, be insisting that in order to contain justified beliefs about the world, a person's belief system would in some sense depend on experience. This gives an essential role to foundations of justification (or knowledge)—grounds of belief that are not true or false and do not themselves admit of justification. It is true that the view would also require coherence among beliefs as an essential element; but a moderate foundationalist could agree that coherence is necessary within a body of justified beliefs such as normal people have, yet insist that this coherence is not a basic source of justification rather than, chiefly, a *product* of the elements, such as grounding in experiential and rational sources, that *are* basic.

If coherentists cannot show that coherence among beliefs is a basic source of justification—as it is far from clear they can—then requiring coherence with experience to make their theory plausible yields a view that is apparently at least compatible with a moderate foundationalism and may well be a version of that view. This may be a welcome conclusion for epistemologists uncommitted on the foundationalism–coherentism issue, but it would be unwelcome to philosophers in the coherentist tradition.<sup>20</sup>

Suppose, however, that moderate foundationalism is correct. We must not suppose that this theory leads easily to an adequate, detailed picture of a typical body of knowledge or justified belief. Moderate foundationalism as so far described—mainly structurally—tells us only what sort of structure a body of knowledge or of justified belief has. It says that if one has any knowledge or justified belief, then one has some direct knowledge or directly justified belief, and any other knowledge or justified belief one has is traceable to those foundations. A belief direct and foundational at one time may be indirect and non-foundational at another; it may gain or lose justification; it may have any kind of content; and some foundational beliefs may be false or unjustified or both.

By leaving this much open, however, moderate foundationalism avoids a narrow account of what is needed for knowledge and justification and allows many routes to their acquisition. For similar reasons, it avoids *dogmatism*, in the sense of an attitude of self-assured certainty, especially concerning claims that are neither self-evident nor obvious. In addition to avoiding this *attitudinal dogmatism*, it rejects, for the same sorts of reasons, at least one version of *epistemological dogmatism*—the version ascribing to us indefeasible justification, epistemic certainty, or the like, where these attributions are unwarranted by our evidence. For moderate foundationalism allows alternative kinds of foundational beliefs for different people and under different circumstances; and, by acknowledging the imperfect reliability of the experiential sources and of many inferences from the beliefs they generate, it also explains why it is so difficult to know that one has knowledge or justified belief, and hence important to be open to the possibility of mistakes.

Moderate foundationalism even allows that a person may not always be able to see the truth of a self-evident proposition. One might, for instance, lack conceptual resources for adequately understanding it. This point should induce humility about how extensive our knowledge is even regarding what is in principle readily known. Ignorance can occur where one would least expect it. Moderate foundationalism also treats reason as a fallible source of belief: we can easily take a false proposition to be true on the basis of a specious sense of its being a priori. This should induce humility about how confident we are entitled to be. Error can occur where it might seem impossible. Foundationalism is committed to unmoved movers; it is not committed to unmovable movers. It leaves open, moreover, just what knowledge is, and even whether there actually is any. These questions must still be faced.

## Notes

- 1 Clearly, there could be devices or strategies by which one can manipulate one's beliefs; what I deny is that we can control belief "at will" (simply by willing it) the way we can normally raise an arm at will. The point is not that the will has no power over belief. For wide-ranging critical discussion of doxastic voluntarism see William P. Alston, 'The Deontological Conception of Epistemic Justification', *Philosophical Perspectives* 2 (1983), 257–99; my 'Doxastic Voluntarism and the Ethics of Belief', *Facta Philosophica* I, 1 (1999), 87–109, reprinted in Matthias Steup (ed.), *Knowledge, Truth, and Duty* (Oxford: Oxford University Press, 2001); and, for critical discussion of Alston's position, Steup's contribution to that collection, 'Epistemic Duty, Evidence, and Internality'. Detailed treatment of our control of our *grounds* for belief and of their relation to our beliefs they support is provided in my 'The Ethics of Belief: Doxastic Self-control and Intellectual Virtue', *Synthese* 161 (2008), 403–18.
- 2 There is dispute about whether people can have infinite sets of beliefs.

I have offered some reasons for doubting this (and cited some of the relevant literature) in 'Dispositional Beliefs and Dispositions to Believe', *Noûs* 28 (1994), 419–34.

- 3 Granted, one could look at the formulation, say by tracing it along a mile-long print-out, and believe that it expresses a truth; but the point is that one could not grasp, and so could not believe, the truth that it expresses. Of course, if we are talking about infinity, the relevant formulations would approach an infinite number of miles in length. For an extensive discussion of the prospects for epistemological infinitism, see Peter D. Klein, 'Human Knowledge and the Infinite Regress of Reasons', *Philosophical Perspectives* 13 (1999), 297–325.
- 4 A highly instructive paper by Peter Klein indicates apparent exceptions. Suppose the trustworthy department secretary told me, last *Thursday* (and knew), that I have an appointment this coming Monday. Now, asked whether I am free at the relevant time Monday, I say 'No, the secretary told me on Friday that I have an appointment then'. Plainly, I can know I have the appointment, though the belief I express now as a basis is false, since I have the wrong day. See 'Useful False Beliefs', in Quentin Smith, *Epistemology: New Essays*, (Oxford: Oxford University Press, 2008). Note three points. (1), on my account of testimony-based knowledge (Chapter 7), I would know that I have the appointment non-inferentially. The false belief that she told me on *Friday* is *offered* as a ground of my belief that I have the appointment on Friday, but the ground on which I know that is her attesting to it. (2) Suppose my statement 'She told me on Friday that I have an appointment Monday' does in a way express my ground. The ground may be that she told me that I have an appointment then, with 'on Friday' functioning like a parenthesis, such as 'and by the way it was on Friday'. Then I do know my ground. (3) What if I have forgotten her testimony, am told that she gave it on Friday, and infer, apparently from this falsehood, that I have the appointment? Now we need a theory. One move is to make a distinction, useful in any case, between the ground's being, as in (2), that she told me that I have an appointment then, and its being, by contrast, something like: It was on Friday that she told me that I have an appointment then (in which the time is important in my thinking). In the latter case I would not know; in the former I presumably would. I thank Claudio de Almeida for introducing this problem to me in a draft ('Knowledge and Benign Falsehoods') written in reference to Klein's paper but long before its publication.
- 5 An item of knowledge can occur in more than one epistemic chain, as when you have two entirely independent sets of premises showing the same conclusion. The regress argument requires one chain, but it allows more than one.
- 6 We may also draw the more general conclusion that if there is any knowledge, there is some direct knowledge. This more general conclusion

follows only on the assumption that if there is any knowledge, then there is at least one knower who has it. This is self-evident for the main sense of 'knowledge'; but if we think of certain books as containing knowledge and then imagine the possibility that all knowers cease to exist while the books live on, it may then seem that there would be (residual) knowledge without there being any knowers (though even here there would have been knowers). Such unpossessed knowledge is discussed in some detail in Chapter 12.

- 7 See Aristotle's *Posterior Analytics*, Books I and II. His argument is importantly different in at least one respect: he spoke of the foundational items as "indemonstrable," which implies that there cannot be any deeper foundations. The regress argument as stated here implies only that one's foundational knowledge is of something that (at the time) one has not demonstrated. This leaves open that one might later demonstrate it by appeal to something "deeper."
- 8 The possibility of combining psychological foundationalism with epistemological coherentism seems quite open to Wilfrid Sellars, widely considered a leading coherentist. See, for example, his 'The Structure of Knowledge', in Hector-Neri Castañeda (ed.), *Action, Knowledge, and Reality: Essays in Honor of Wilfrid Sellars* (Indianapolis: Bobbs-Merrill, 1975).
- 9 This case is styled after one of Klein's; see esp. 'Useful Falsehoods', 36–40, cited above. His treatment of it differs from the one suggested here but may be consistent with that.
- 10 The point here does not presuppose the mistaken idea that coherence is *closed under negation*, in the sense that if a set of propositions is coherent, so is the set whose members are their negations. Consider the set consisting of *p*, *q*, and *p explains q*. Negating these need not yield a coherent set. This is one way coherence differs from mere consistency.
- 11 This is the position taken by Laurence Bonjour in *The Structure of Empirical Knowledge* (Cambridge, MA: Harvard University Press, 1985). It should perhaps be noted that he has rejected the coherentist epistemology of this book in, e.g., 'The Dialectic of Foundationalism and Coherentism', in John Greco and Ernest Sosa (eds.), *The Blackwell Guide to Epistemology* (Oxford: Blackwell, 1999).
- 12 Recall the wheel described in discussing circular causation above. Does the fact that the topmost eastern section is in place not explain why the topmost western section, which is contiguous with it, is in place, and isn't the converse also true? Only, I suspect, if this comes to saying that given these facts we can infer that each is in place. Why each is in place is explained by the same thing: the overall pattern of forces including the support provided by the ground. Each is in place because the gravitational force pulling it backward and downward is matched by a gravitational force pulling it forward and holding it up: both phenomena are indeed explained by the "same thing"—the qualitatively identical

- forces—but not by the same thing in the sense of the other, qualitatively identical phenomenon. Explanation by two phenomena that are “exactly alike” exhibits a kind of mutuality, but it is not the same as explanation of each of two exactly similar phenomena in terms of the other.
- 13 As suggested in our discussion of self-evidence, there are some beliefs, say that if  $x = y$  then  $y = x$ , and that I exist, which it is plausible to think cannot be unjustifiedly held.
  - 14 On the topic of practices of justification, Ludwig Wittgenstein’s *On Certainty* (Oxford: Basil Blackwell, 1969) is a valuable source. He is often cited as stressing that there comes a point at which one says: *My spade is turned* (a foundationalist metaphor).
  - 15 This line of thought is suggested by what Laurence Bonjour, in *The Structure of Empirical Knowledge*, calls “the observation requirement.” For extensive discussion of his theory there and of coherentism in general, especially that of Keith Lehrer, see John W. Bender (ed.), *The Current State of the Coherence Theory* (Dordrecht: Kluwer, 1989).
  - 16 A slightly different formulation may be required if, for the sorts of reasons to be given in Chapter 10, knowledge does not entail justification; but the formulation given will serve here. Here and elsewhere the reference to foundational beliefs is to those that are justified; I also omit an other-things-equal clause appropriate after the ‘if’ in clause (3). For a highly detailed statement of a moderate foundationalism, see Paul K. Moser, *Knowledge and Evidence* (Cambridge: Cambridge University Press, 1989).
  - 17 The independence principle cited here is not the only one that seems sound. For instance, it is plausible to hold that one’s justification for a proposition also rises (other things being equal) the larger the number of factors one is appropriately aware of that do support it, whether or not one believes them to do so. If the independence principle were weakened by eliminating the requirement that one has some justification for taking the relevant sources to support the belief in question, we may certainly say that more justification is conferred (other things equal) by factors justifiedly taken to support the belief than by those unjustifiedly taken to do so.
  - 18 It is a strong foundationalism, especially the kind found in Descartes’ *Meditations*, that is influentially criticized by Richard Rorty in *Philosophy and the Mirror of Nature* (Princeton, NJ: Princeton University Press, 1979). Many of Rorty’s criticisms do not hold for the moderate foundationalism developed in this chapter. His doubts about the very idea that the mind is a “mirror of nature,” however, may cut against at least the majority of plausible epistemological theories, depending on how much is built into the metaphor of a mirror. This book as a whole can be seen as a case for a version of realist epistemology, and some aspects of Rorty’s challenge are treated at least implicitly in Chapters 12–14 and in parts of other chapters, such as the sections on phenomenalism and truth.

- 19 Keith Lehrer provided an influential statement of this view in *Knowledge* (Oxford: The Clarendon Press, 1974): having said that “complete justification is a matter of coherence *within* a system of beliefs” (p. 17, emphasis mine), he added, “There is no exit from the circle of ones [sic] own beliefs from which one can sally forth to find some exquisite tool to measure the merits of what lies within the circle of subjectivity” (pp. 17–18). Such sensory states as an impression of green grass are among the excluded tools. Further indications of why a coherentist view disallows appeal to experiential and other non-truth-valued states as justificatory are given by Wilfrid Sellars, ‘The Structure of Knowledge’, in Castañeda, *Action, Knowledge, and Reality*, cited above; and Donald Davidson, ‘A Coherence Theory of Truth and Knowledge’, in Dieter Hendrich (ed.), *Kant oder Hegel* (Stuttgart: Klett-Cotta, 1983). It should be noted that Davidson has written an afterword to this paper. Here he says, regarding “The main thrust of ‘A Coherence Theory’,” that “the important thesis for which I argue is that belief is intrinsically veridical.” See ‘Afterthoughts, 1987’, in Sven Bernecker and Fred Dretske (eds.), *Knowledge: Readings in Contemporary Epistemology* (Oxford: Oxford University Press, 2000), p. 427. This view bears a foundationalist interpretation: beliefs, being intrinsically veridical, are as such *prima facie* justified, even if weakly; hence they are defeasibly foundational. Incoherence would defeat them; but particularly if, as Davidson says here, “Coherence is nothing but consistency” (p. 427), it would not be plausible to take coherence to be a basic source of justification.
- 20 The idea of enriching coherentism by making coherence with experience an essential element in coherentist justification is proposed and defended by Jonathan L. Kvanvig and Wayne D. Rigg, ‘Can a Coherence Theory Appeal to Appearance States?’, *Philosophical Studies* 67 (1992), 197–217. This paper deserves study. Here I raise just one difficulty. Although they grant that “coherentism arises historically because of dissatisfaction with the foundationalists’ picture” (p. 199), they characterize a foundationalist warranting relation in a way that does not distinguish it from the relation coherentists take to confer justification.

One such account could claim that a belief is foundationally warranting just in case the evidence for it is an appearance state involving the same content as that of the belief. For example . . . perhaps my belief that something is red is intrinsically warranting because it appears to me that something is red.

(p. 199)

A foundationalist need not take an appearance state, such as a sensory impression of red, to have the same content as a belief (at least a propositional belief, the kind apparently in question here): a propositional content in virtue of which the belief is true or false, e.g. “that something

is red.” Rather, the experiential content is qualitative, roughly a matter of what properties one is sensorily acquainted with in having the experience; that content may *correspond* to certain propositions but is not itself truth-valued. Such a content might be an appearance of red but not the proposition that “it appears to me that something is red.” The latter is a candidate to enter into a coherence relation with beliefs or their contents. Perhaps Kvanvig and Rigg are thinking of experiential justification of belief as possibly working through beliefs or other states which have propositional content and truth value; this could explain why they find such justification available to coherentism. If, however, experiential justification could work that way, then one could still have a coherent system of beliefs (perhaps even objectual beliefs) that goes against experience. Beliefs about one’s states—such as the (appearance) “belief that something is red”—would have to play a role, but those states would not be any kind of bedrock grounding these beliefs, even if the beliefs happened to be based on them. The problem, then, is that either the coherence-with-experience approach assimilates coherentism to a kind of foundationalism or it fails to capture the role of experience, which seems essential for a body of justified beliefs about the world.





## **Part Three**

The nature and scope of  
justification and knowledge



# 10 The analysis of knowledge

Justification, certainty, and reliability

- **Knowledge and justified true belief**
- **Knowledge conceived as the right kind of justified true belief**
  - Dependence on falsehood as an epistemic defeater of justification*
  - Knowledge and certainty*
  - Knowing and knowing for certain*
  - Knowing and making certain*
- **Naturalistic accounts of the concept of knowledge**
  - Knowledge as appropriately caused true belief*
  - Knowledge as reliably grounded true belief*
  - Reliable grounding and a priori knowledge*
- **Problems for reliability theories**
  - The specification problem*
  - Reliability and defeat*
  - Reliability, relevant alternatives, and luck*
  - Relevant alternatives and epistemological contextualism*

## 10 The analysis of knowledge

### Justification, certainty, and reliability

Knowledge arises in experience. It emerges from reflection. It develops through inference. It has a distinctive structure. The same holds for justified belief. But what exactly is knowledge? If it arises and develops in the way I have described, then knowing that something is so is at least believing that it is. But clearly it is much more. A false belief is not knowledge. A belief based on a lucky guess is not knowledge either, even if it is true.

Can something be added to the notion of true belief to yield an analysis of what (propositional) knowledge is, that is, to provide a kind of account of what constitutes knowledge? Plato addressed a similar question. He formulated an account of knowledge (though in the end he did not endorse it) which has sometimes been loosely interpreted as taking knowledge to be justified true belief.<sup>1</sup>

For Plato, 'belief' would represent a grade of cognition lower than knowledge. But if we substitute, as most interpreters of Plato would have us do, some related term for 'belief', say 'understanding', then the account may be nearer to what Plato held and would be closer to some of the historically influential conceptions of knowledge. In any case, the notion of belief is wide and subtle; and one or another form of the justified true belief account prevailed during much of this century until the 1960s.<sup>2</sup> What can be said for it?

### Knowledge and justified true belief

What is not true is not known. When we claim we know something and later discover that it is false, we sometimes say things like 'Well, I certainly believed it'; but we do not seriously maintain that we knew it. One might say 'I just *knew* it', but this is usually taken to exhibit an inverted commas use of 'know', a use in which 'know' stands in for something like 'was certain'. If we seriously claimed we knew it, others would likely conclude that (for instance) we do not really believe that it is false, or perhaps are using 'I knew' to mean 'I felt great confidence', as in 'I just knew I'd win—I still can't really believe I lost'. In cases like the commonsense ones just described, when truth is subtracted from what appears to be knowledge, what remains is not knowledge but belief.

These points suggest that knowledge is at least true belief. Admittedly, people who feel certain of something, for instance that a friend is angry, may say that they do not believe it, but know it. This is best understood, however, to mean that they do not *merely* believe it, but know it.

Similarly, it may be misleading to say "I believe he's angry" when I think I also know it—unless I intend, for instance, to indicate caution or perhaps polite disagreement. But it is often misleading to say less than one is fully entitled to say.<sup>3</sup> My saying that I believe he is angry may be misleading precisely because I am expressing only part of what I am fully entitled to express: that I know he is. For I am thereby suggesting that I do not know, or perhaps even doubt, that he is. If this point is what explains why my statement is misleading, that confirms that knowing implies believing.

Does knowing something also imply justifiedly believing it? If it does, that would explain why a true belief based on a lucky guess is not knowledge. If, from a distance, I see Jim walk hurriedly down the hall and simply guess that he is angry, I am not justified in believing that he is angry. If my belief turns out to be true, it still does not constitute knowledge, and its lacking justification apparently explains why not. Now suppose I go by his office and see him briskly shuffling papers and angrily mumbling curses. At this point I might come to know that he is angry; and my acquiring knowledge that he is can be explained by my having acquired evidence which justifies my true belief that he is.

Still, could a true belief that is not justified constitute knowledge? Suppose I simply see Jim briskly shuffling papers as I pass his office, but do not hear any curses. A bit later, I see him walk hurriedly down the hall. Given that I know his fiery temperament, I might have just enough evidence to give me some reason to believe he is angry, even though I am not quite justified in believing this. Might I now have a kind of low-grade knowledge that he is angry? This is doubtful. My evidence for believing this is not strong. But the case does show this much: that as our evidence for a true belief mounts up in a way that brings us closer to justification for holding it, we also tend to get closer to knowledge. These and similar points support the view that justified belief is an element in knowledge. This view is highly plausible, and—for now—I want to assume it.

We are, then, on the way toward an analysis of knowledge. For it looks as if we have a very substantive threefold necessary condition for (propositional) knowledge: it seems that knowledge is at least justified true belief—that we know something only if we believe it, it is true, and our belief of it is justified. Still, a correct, illuminating analysis, one that provides a good account of the nature of what is being analyzed, must also provide sufficient conditions. It might be true that I know something only if I justifiedly and truly believe it, yet false that if I justifiedly and truly believe something, I know it.

It apparently is false that if we have a justified true belief, we (always) have knowledge. Suppose that when I first visit the Wallaces I have no idea that they have a photographic collection which includes realistic, life-size pictures

of themselves. When I approach the doorway to their living room down a long hallway, I see, just twelve feet before me, and constituting all I can see through the doorway, a life-size picture of Jane, standing facing me and smiling like the good hostess she is, with the background in the picture looking just like the living room's rear wall. I say 'hello' before I get close enough to realize that I see only a photograph of her taken against the background of that very wall. I discover that the picture is so lifelike that this happens to everyone who knows Jane and enters unaware of the photograph. I might thus be quite justified, momentarily, in my belief that Jane is opposite me. As it happens, however, Jane is standing opposite me—in the next room, right behind the wall on which the picture is hung. My belief that she is opposite me is thus true, as well as justified. But I do not know that she is opposite me.<sup>4</sup>

This example shows that if we analyze knowledge as justified true belief, our analysis is too broad. How might we improve it? If taking justification to be crucial is on the right track, we can restrict the kind or degree of justification involved. We might, however, suspect that justification is not central after all, but only correlated with something that is. We might then seek an account of knowledge in which justification is not central to understanding knowledge. There are many approaches of both kinds. I want to consider two of each, starting with the "justificationist" accounts.

### **Knowledge conceived as the right kind of justified true belief**

In the photographic case, something seems wrong with the kind of justification I have. It is sometimes said to be defeated—alternatively, *defective*—where this is not to say that it is undermined or overridden, as in the more common cases of defeated justification we have so far noted, but rather (in part) that it is prevented from playing what seems to be its normal role in such a case, namely, rendering a true belief knowledge. Contrast this kind of defeat of justification with the common kind that undermines or overrides justification—as when one discovers a witness one had believed was lying and is thus no longer justified in believing the testimony. Call the former epistemic defeat: it eliminates the power of the justification to turn a true belief that acquires that justification into knowledge. In that sense, it vitiates the justification, eliminating its characteristic power to raise the status of a merely true belief to that of knowledge. Perhaps, then, with epistemic as opposed to justificational defeat in mind, knowledge might be analyzed as *undefeatedly justified true belief*.<sup>5</sup> This idea is worth exploring.

### ***Dependence on falsehood as an epistemic defeater of justification***

How is (epistemic) defeat to be characterized? One natural view is that the justification of a belief is defeated provided the belief depends on a falsehood.

A dependence on falsehood is a bad thing from the point of view of the candidacy of a belief to constitute knowledge. This is in part because, even when a belief which depends on falsehood is true, that may be just by good luck. It is good luck that Jane happens to be standing straight in front of me.<sup>6</sup> But as our lifelike photograph example shows, a belief true just by good luck does not constitute knowledge.

There are at least two ways in which a belief might depend on falsehood. First, it might depend on a falsehood in the sense that it would not be justified except on the basis of one's being (situationally) justified in believing a falsehood about the subject in question (say, Jane). This is a kind of justificational dependence (dependence for justification), which I will call *presuppositional dependence*. In the photographic case, my belief that Jane is opposite me depends presuppositionally on the falsehood that I am seeing her directly (or at least in a way that does not misrepresent her location).

The point is not that in order to know she is opposite me I would have to believe the false proposition that I am seeing her directly; rather, my belief that she is opposite me *epistemically depends* (depends for its claim to be knowledge) on this proposition. Not only does it seem to be because this proposition is false that I do not know Jane is opposite me; it is also the kind of proposition whose truth is central for grounding my would-be knowledge and whose falsity I would tend to be surprised to discover.

The second case of dependence on falsehood is *psychological dependence*: a belief might psychologically depend on a falsehood in the causal sense that one has the belief by virtue of holding it on the basis of believing a falsehood. In this kind of case my would-be knowledge is sustained by a false belief, which is a kind of inadequate foundation. My belief about Jane would psychologically depend on falsehood if, say, I knew about the Wallaces' life-size photographs, yet trusted my vision and believed that Jane was opposite me on the basis of concluding that this time I was viewing her directly. I am not viewing her directly, so my underlying belief is false.

Unfortunately, the appeal to a false presupposition, or even to other kinds of dependence on falsehood, does not always explain why a justified true belief is epistemically defeated and fails to constitute knowledge. Recall the lottery with a million coupons. You might have a justified true belief that you will lose, but you do not know that you will. You might possibly win. What falsehood defeats your justification here? You are not making any mistake, but simply do not have the right *kind* of positive ground for knowledge.

It might seem that your belief that you will lose the lottery depends on the false proposition that the outcome of a chance process can be known beforehand by merely calculating odds. But does your belief depend on this? You might reject this and still believe—even justifiedly—that you will lose, whereas I could not reject the false presupposition that I see Jane directly and still believe (justifiedly, at least) that she is in front of me.

We should not hold, then, that in the lottery example either your belief or its justification depends on the falsehood about foreknowledge of chance outcomes. Points like these do not show that no version of the undefeatedly



justified true belief analysis of knowledge will work. One might, for instance, try to explain why justification is epistemically defeated in the lottery case even though it is based on as high a probability as one likes. This effort may lead in the direction of the next justificationist account of knowledge I want to consider.

### *Knowledge and certainty*

The lottery example suggests that knowledge requires one's possessing *conclusively* justified true belief, belief justified in such a way that its truth is guaranteed by what justifies it. For we may plausibly claim that if the evidence guarantees that you will lose, say because it includes knowledge of the lottery being fixed in favor of someone else, then you would know you will lose. Moreover, *conclusive* justification is presumably not liable to defeat (a point that a defeasibility view of knowledge can also make use of).

Different theories offer different accounts of a guarantee of truth (as will be apparent when we discuss skepticism in Chapter 13). The lottery example supports the view that the right kind of guarantee is not simply a matter of high probability. After all, we can have as many tickets as we like and you would still not know yours will lose. Thus, in this kind of case, no matter how probable it is that you will lose, your justification is not sufficient for knowledge.<sup>7</sup> This approach to conclusive justification would not entail that extremely high probability could *never* suffice for knowledge, say when it represents the likelihood of there being print before us given our present experience of reading this; but in that kind of case there is apparently no randomizing process whose outcome determines whether our beliefs are true. That kind of randomness, then, might be said to rule out conclusive justification. There are other disanalogies between the perceptual and lottery cases (more than can be even noted here). As I look at this print, for instance, I may assume that there is not (and perhaps never has been) someone in exactly the same evidential position—assuming we can determine that—who is mistaken in an exactly similar belief that there is print.

Another reason to think that knowledge requires conclusive justification is that knowing is often closely associated with certainty. When I wonder if I know, I may ask myself how I can be certain. I also sometimes wonder if what I believe is certain. Particularly in the latter case, I am thinking of the status of the proposition in question, not of *psychological certainty*, which is, roughly, great confidence of the truth of what one believes. If I am confident enough that (some proposition) *p* is true, I am (psychologically) certain that it is and certain *of* it; and if I am certain of it, I am supremely confident that it is so and certain *that* it is so.

Given that there are two kinds of certainty, it is important to see that the question 'How can I be certain?' does not concern only psychological certainty. It typically means something like 'How may I justifiedly be

(psychologically) certain?’ and there it may be intended to ask whether propositional certainty is possible. If, however, I say, not that I am not certain but that *it* is not certain, that your ticket will lose, I am referring to propositional certainty, roughly, the certainty a proposition has when there are extremely strong grounds for it, grounds that guarantee its truth.

I want to leave two things open: first, what kinds of grounds guarantee truth; second, how readily available the grounds of propositional certainty must be, if readily available at all, for instance whether ordinary reflection on what evidence one has would reveal them. Saying that something is certain surely implies that one thinks sufficient grounds are in some sense available, even if only by a careful study of the matter and perhaps consulting others about it or otherwise getting new information. But for a proposition just to be certain, the kind of availability (if any) is more difficult to assess.

Given these connections between knowledge and certainty, one might hold that knowledge is constituted by conclusively justified true belief, meaning that (1) the believer may justifiably be psychologically certain of the true proposition in question and (2) this proposition is so well-grounded as to be itself propositionally certain. Knowledge constituted by such a justified belief may be (and has been) considered a case of *epistemic certainty*.<sup>8</sup>

An analysis of knowledge as constituted by a belief exhibiting epistemic certainty seems too narrow. It would, for instance, apparently rule out most knowledge based on testimony. If Jane tells me that she wants to meet to discuss something, and I know her well and have no good reason to doubt her word, may I not know that she wants to meet with me? Yet I do not have conclusive justification, nor does her testimony render it certain that she wants to meet with me. Unlikely though it is, error is barely possible; she could act out of character and deceive me (or herself).

### *Knowing and knowing for certain*

Does knowing imply, if not conclusive justification of the belief constituting knowledge, then at least the certainty of the proposition known (epistemic certainty)? In the case described, I doubt that it is propositionally certain that Jane wants to meet with me (but the notion of propositional certainty is vague, and it is often difficult to tell whether it applies). My knowledge here is apparently not knowledge of something that is certain.

Indeed, we sometimes speak of knowing something for certain, implying a contrast with simply knowing. Imagine that Tom tells me that Emma has left town, but I believe him to be mistaken and say so. Someone probing my grounds for saying Tom is mistaken might ask if I know this for certain. This might be asked not from doubt about whether I know, but to find out if the proposition that he is mistaken is certain, perhaps because much hangs on it, as in a criminal trial, where I must testify. It might also be asked from a desire to determine what kind of basis I have for my claim. An answer would

be that I just took her from the station to her hotel and I can't imagine her having already left when she's on the conference program. The existence of such cases suggests that what is not known for certain still can be known.

It is interesting to compare knowing and knowing for certain with simply knowing someone and knowing the person *for* a practical joker. We can know a person who is a practical joker without knowing, or in any way taking, the person to be a joker. If the parallel holds, it suggests we can know that *p* without knowing it for certain.

Perhaps, however, what can be known at all can always be known for certain, as I might come to know for certain—provided I do enough checking into her motivation—that Jane wants to meet with me. But even if what is knowable can be known for certain, it is doubtful (as examples to be given will also suggest) that everything that is known is certain. Still, is it even true that whatever is known must at least be such that it can be certain? Our example suggests that knowledge need not meet this standard: I might know that Jane wanted to meet even if she has just died and there are no additional evidences—such as recollections by third parties, letters by her—on the basis of which this can be certain.

One might reply that knowledge is always of the *sort* of proposition that can be certain. But consider propositions about the past, such as that a ship sank in a certain lonely region in the Pacific Ocean. Perhaps these may be knowable, yet not even be the sort of thing that can (for human beings, at least) be certain (the evidence of its apparent traces at the bottom is good enough for knowledge, but no further evidence, such as eye-witness accounts, can be obtained regarding where it sank).

### *Knowing and making certain*

Perhaps, however, these points show only that I cannot *make* certain that Jane wants to meet with me. There might still be a basis for this proposition which renders it (propositionally) certain. I will continue to leave open the question of whether what is known is the sort of thing that can be certain, since knowledge, not certainty, is my main concern here. But it will help, in that connection, to explore how an understanding of the notion of making certain may bear on the view that knowledge requires conclusive justification.

If, when we already know something, we can make certain that it is so, then there is reason to think that conclusive justification is not required for knowledge. Suppose I lock the back door and, as I drive off, clearly recall doing so. Still, if someone asks me if I am sure I did, I may truly believe I know I did, yet still check to make certain I did. Now when we need to (or even can) make certain of something we know, it would appear that it need not *be* either certain or conclusively justified for us. Getting conclusive justification seems to be the main point of making certain, though on some views the latter may be weaker, in that there may be cases in which we make certain of something but still lack utterly conclusive justification for it.

It might be replied that in the case in which making certain consists of getting further evidence, 'make certain' means not 'make *it* certain' but, roughly, 'make sure it *is* certain', and that if I really knew it, it was certain in the first place. Let us assume for the sake of argument that 'make certain' means 'make sure it is certain'. Now suppose I do not make certain that I locked the door, because it begins to hail and I must leave before I can check the door. This does not show that I do not know I locked the door; and, on later finding that it was locked, I could be correct in saying that I was right all along to think I knew. So, even if making certain were a matter of making sure the proposition is certain, it does not seem either that one needs to do this in order to know the proposition, or even that the proposition must *be* certain, in order to be known.

Moreover, supposing I did know all along that the door was locked, it does not follow that this was certain all along. I had good reason, at least at the time when I could not check, to think it was not certain. All things considered, the possibility of making certain of what we already know suggests that knowing a proposition does not entail its being certain. Further, if, as it seems, we can know something, yet make certain it is so, then apparently we can also know it without being *conclusively* justified in believing it.

Notice that similar points apply to what we know from memory. Even on topics with respect to which our memory is highly reliable, the justification our memory beliefs have is generally not conclusive. Even if I can recite a stanza from memory, my justification for believing I have it right need not be conclusive. Yet I may well know that I have it right, and confirm that I do when I look it up to make certain I do and I find that it reads just as I thought.

## Naturalistic accounts of the concept of knowledge

Perhaps we should consider a quite different approach. Must we appeal to the notion of justification to understand knowledge? Suppose we think of knowing as registering truth, somewhat as a thermometer registers temperature. Knowledge, so conceived, results from the successful functioning of our epistemic equipment, which consists above all of finely tuned perceptual, memorial, introspective, and rational instruments.

The thermometer analogy holds not just for propositional knowing but also both for simple knowing—knowing, *by acquaintance as opposed to description*, persons and objects—and for objectual knowing, which is knowing something *to be* a certain way. But propositional knowing is my main topic here; and, from a study of the chapters on perception, one could largely adapt to the other cases what emerges about propositional knowledge.<sup>9</sup>

The view that knowledge consists in suitably registering truth goes well with the idea that we are biological creatures with sense receptors that gather information and with mental capacities that integrate it. Perhaps, then, knowledge can be analyzed *naturalistically*, that is, using only the kinds of concepts the sciences, especially the natural sciences, use in understanding

things. These prominently include “observation concepts,” such as those of color and shape, height and weight, number and motion. A naturalistic account appeals not to normative notions—“value-laden” notions, in one terminology—like that of justification, but (largely) to physical, chemical, biological, and psychological properties, together with causal relations among these.

I want to consider two naturalistic approaches. The first emphasizes the role of causation in producing our knowledge, as with perceptual beliefs caused by the perceived object. The second approach stresses the reliability of the processes, such as seeing, through which knowledge arises.

### *Knowledge as appropriately caused true belief*

On the causal theory, knowledge is true belief caused by something connected with its truth in a way that makes it plausible to call the belief knowledge. Roughly, knowledge is appropriately caused true belief, in which appropriate (causal) production of a belief is production of it in which the fact, object, event, or other thing in virtue of which the belief is *true* plays a certain role in generating or sustaining the belief.

In the examples of knowledge that best support the view, the belief in question is apparently a case of knowledge because it is caused in a way that guarantees its truth. Thus, I know that there is a green field before me because the field itself plays a major part, through my vision, in causing me to believe there is a green field before me. I know that Jane wants to meet with me because her wanting to do so plays a major part in causing her to say she does, and thereby in causing me to believe that she does. I know that the stanza I recite from memory has four lines because its having them is a major causal factor, operating through my memory, in my believing that it does.

The causal view can even accommodate knowledge of the future. I know that I am going to continue thinking about knowledge for a long time. That truth (about the future) does not cause me to believe this; but that truth is causally connected with my belief, and in a way that suggests why the belief may be expected to be true. For what causally explains *both* why the proposition I believe about the future is true and why I believe it is the same element: my intending to continue thinking about knowledge. Here my future-directed belief is knowledge, but not by virtue of being produced by the thing it is about—my future thinking—for that has not occurred.

Does this view of knowledge of the future show that as the relevant facts lie in the future, knowledge need not represent “the facts,” as the common-sense view has it? The case need not be so interpreted. *Representing* facts does not require being *caused* by them. Recall my belief about the future. This belief constitutes knowledge, but not because what the belief is about is the way the belief represents it *and* causes the belief, as when the green field’s being before me causes me to believe that it is before me. Rather, such a future-directed belief correctly represents what it is about in part because *the belief itself* causes that state of affairs.

Moreover, the causal theory is right about this much: my belief that I will continue thinking about knowledge is caused by something—my intention to continue thinking about it—of a kind that makes it at least likely that I will be as the belief represents me. Roughly, not only can knowledge be produced by things that are known, as in perceptual cases, so that knowledge is related to what is known as *effect to cause*; both knowledge of certain facts and the facts known can also be *common effects of the same causes*, as in the case of knowledge of the future.

There are, however, serious troubles for the theory that knowledge is appropriately caused true belief. One problem is how to apply the basic idea—that what underlies the truth in question is a causal factor in the grounding of the belief of that truth—to a priori knowledge. How might what underlies the truth that if one tree is taller than another then the second is shorter than the first be causally connected with my believing this truth? This truth is not (in general) perceptually known, nor is its status dependent on any particular object in the world, as is the case with the (empirical) knowledge to which the causal theory best applies.

It may be that the only way a truth can be causally connected with a belief so as to render it knowledge is through a connection with something in the world that does at least partly cause (or is at least partly an effect of) the belief. The truth that there is a green field before me is about an object that produces visual impressions in me. But the strictly a priori knowledge just cited does not depend on trees in that way. It does not even depend on there ever being any trees. It seems to be based simply on a grasp of the concepts involved, above all that of a tree and that of height. My having this grasp does not appear to imply causally interacting with those concepts (supposing it is even possible to interact causally with concepts).<sup>10</sup> This is not to say that the belief has no causal ground, such as the understanding or the comprehending consideration of the relevant proposition. The problem is that a belief constituting a priori knowledge that *p* seems to lack the kind of ground the causal theory requires: the a priori fact that *p* does not cause it, nor are this belief and that fact common effects of the same causes.

### *Knowledge as reliably grounded true belief*

There is another serious problem concerning the causal account, this time in relation to empirical beliefs. When we understand it, we can see the rationale for a different way of understanding knowledge. As in many instances, the trouble arises from examples of justified true beliefs that do not constitute knowledge.

Consider a case in which something causes me to have a true belief, yet that belief is not knowledge. Suppose Tom tells me, on the basis of his knowing it, that Jim is angry, and as a result of his testimony I believe this. My belief might be justified and true. But imagine that, although I have no reason whatever to believe this about Tom, he is in general highly unreliable, and sometimes lies, in what he says about Jim. The mere fact of Tom's

unreliability precludes my knowing through his testimony that Jim is angry. Even if Tom knows Jim is angry, and knows it because he observes Jim acting angrily, his knowledge is not transmitted to me. For he might well have said this even if Jim had merely acted, say, hurriedly, and was not angry. (This shows that even if, as held in Chapter 7, testimony transmits knowledge to a belief based on it *only* if the attester knows the proposition in question, it need not do so even when the attester *does* know it.)

It is important to see that although the causal connections here seem to be what they usually are in testimony cases, I do not acquire knowledge from Tom's testimony. Jim's anger causes Tom to believe him angry; Tom's belief (partly) causes his telling me that Jim is angry; his telling me this causes me to believe it. But, though I have a justified true belief that Jim is angry, I do not know it. For while Tom has it right this time, he is in general unreliable regarding Jim.

The testimony example brings out something very revealing. It suggests that the reason I do not know on the basis of Tom's testimony is that it is not reliable. By contrast, perception normally is reliable: normally, at least, we may justifiably count on the beliefs it typically produces, such as beliefs based on touch or vision, to be true. We may also presume that perception is also reliable in the sense that the vast majority of beliefs it produces are in fact true. When there is a photograph that we are unaware of, however, what we apparently see through it is typically not a reliable indication that this very thing is before us. Cases of these sorts suggest that we might plausibly analyze knowledge as *reliably grounded true belief*.<sup>11</sup>

### *Reliable grounding and a priori knowledge*

To see how this approach works, recall Tom's testimony about Jim. Suppose that Tom is only very occasionally mistaken about Jim. Might I then acquire knowledge on the basis of Tom's testimony? A crucial question is *how* reliable a belief-producing process, such as testimony, must be to yield knowledge. The theory gives us no precise way to answer this.

The theory can be defended on this point, however, by noting that the concept of knowledge is itself not precise. Thus, there may be times when, no matter how much information we have, we cannot be sure whether someone knows or not, just as, because the term 'bald' is vague, we cannot always be sure whether it applies, no matter how much information we have (including the number of hairs on the person's head). It might be added that as the reliability of Jim's testimony goes up, so does our inclination to say that I know on the basis of it. This proportionality seems to confirm the reliability theory.

Even a priori knowledge might well be accommodated on this view. For it is at least normally produced by grasping concepts and their relations, or by certain simple valid inference on the basis of beliefs grounded in such a grasp; and these processes of producing belief seem reliable. Consider the a priori



proposition that every sphere encloses a space. Suppose that understanding of abstract entities and their relations—say of spheres, spaces, and the inclusion relation between them—entails a kind of direct contact with these entities, a kind such that the entities form an essential part of the very content of that understanding. Then they play an indispensable if indirect role in sustaining beliefs which, on the basis of this understanding, are justified *a priori* or constitute *a priori* knowledge. These beliefs, then, have understanding as a causal sustaining ground; the relevant understanding in a sense contains the abstract facts that ground the truths known; and by virtue of this apprehensional causal grounding of the beliefs, they are *reliably* based on the facts they represent. It appears, then, that in both the empirical and *a priori* cases, when we know, we have reliably registered the truth.

### Problems for reliability theories

The reliability theory apparently does receive support from the kind of correlation illustrated above: the tendency to count my true belief about Jim as knowledge apparently varies with the tendency to regard the belief's testimonial basis as reliable. But perhaps our underlying thought in so speaking about the belief is that the more reliable Tom is, the better is my justification for believing what he says. If so, then the reliability theory might give the right results here because it draws on the role of justification as a constituent in knowledge.

To be sure, neither reliabilists nor their justificationist critics need hold that I must *believe* anything specific about Tom's reliability in order to acquire justified beliefs from his testimony. But it might be argued that my knowledge has a presuppositional dependence on the proposition that he is sufficiently reliable to *justify* my accepting his testimony, and that it is either because this presupposition is false, or because I lack justification for believing it, that my justification for believing his testimony is defeated in the first place. Thus, it might be argued that even if the reliability account is correct about the conditions a belief must meet to constitute knowledge, its success may be due to its tacit dependence on the justificationist concepts it seeks to abandon.

### *The specification problem*

There is a different kind of problem that must also be faced by the reliability theory. This difficulty seems deeper than the question of how reliable a process has to be in order to ground knowledge. It concerns how to specify what is reliable in the first place. It will not do to say, for instance, simply that the reliable processes we are talking about are mainly those by which the experiential and rational sources of knowledge produce belief. This is not obviously wrong, but it leaves too much undetermined.

Consider vision. Its reliability varies so much with *conditions of observation*



that it would be wrong to say without qualification that it is a reliable belief-producing process. It might seem that we may say this. It is reliable in producing beliefs in good light with the object of vision near enough relative to the visual powers of the perceiver. But this claim will not do without qualification either. It does not rule out external interferences like deceptive photographs, such as the one of Jane. It also fails to rule out internal interferences such as hallucinogenic drugs. These interferences might produce false beliefs about objects that we in fact see and about which we also have many true beliefs, as when, after brain damage, we hallucinate a dark blight on a green tree which we otherwise see plainly as it is.

There are, moreover, so many possible factors that affect reliability that it is not clear that we can list them all without using blanket terms such as 'too far away' as applied to the object, and 'insufficiently attentive' or 'not acute enough' as applied to the perceiver. These terms are not only quite vague; the more important point is that they may be argued to come to something like 'too far to be reliably (or justifiedly) judged', 'too inattentive to form reliable (or justified) beliefs', and 'not acute enough for reliable (or justified) judgment of the features of the object'. If so, their interpretation may well depend on our already having a good philosophical understanding of reliability (or justification), and they are thus unlikely to help us much in clarifying reliability; or if they do, it is because we are relying on a different theory.

Suppose we can devise a vocabulary that overcomes these problems. A related difficulty may persist. Belief production might be reliably described in one way and unreliably described in another. Hence, even if we are able to specify what, in general, a reliable belief-producing process is, we need a way of deciding what reliable-process description to use in order to understand a particular case. Recall my seeing Jane in the photograph and thereby believing that she is opposite me. Suppose we say—what seems correct—that my belief arises from a process of seeing someone in a photograph that (at the time and in the physical in question) accurately shows the person's features and general location. If this kind of basis suffices for knowledge, then my belief presumably should constitute knowledge. For the picture shows her to be where she is: opposite me.

Suppose, on the other hand, we say something else that applies to the grounding of my belief that Jane is opposite me: that the belief-producing process is one of seeing a woman in a picture which gives the false impression that she is *directly* in front of one. Then my belief arising from the process is clearly not reliably produced—since usually in such cases the person is not opposite one at all—and the belief should thus not be knowledge. The trouble is that both descriptions apply to the production of my belief. Using one description, the theory apparently implies that I know; using the other, it implies that I do not.

How can the theory enable us to choose between the two correct reliable-process descriptions, or justify our choosing whatever kind of description it accepts? Call this the *specification problem* (or description problem; it is also

called the *generality problem*, since a major issue is how general the descriptive terms should be). If we first have to decide whether I know that Jane is in front of me by relying on some quite different understanding of knowledge and only in that light can we frame a description, the theory would seem to give us very limited help in understanding knowledge. For the theory itself can apparently be put to work only insofar as, in the light of some alternative account of knowledge, we already understand knowledge at least well enough to be in a position to tell systematically, for a vast range of true beliefs, whether or not they constitutes knowledge.<sup>12</sup>

This point, however, might be said to be at best exaggerated and to show no more than that to use the reliability view we need a good intuitive grasp of the concept of knowledge. That is plausible. Let us accept it for the sake of argument.

The deeper point is that if we seek to clarify knowledge (or justification) by appeal to reliable belief-grounding processes naturalistically understood, then we need a way of explaining what those processes are without inadmissibly appealing, in our explanation, to the concept of knowledge (or justification). A belief that is knowledge should be such because it is reliably grounded true belief; a reliable belief-grounding process should not be characterized as the kind that yields, say, perceptual knowledge.<sup>13</sup>

Similarly, if we have to find the right reliable-process description in terms of what I am *justified* in presupposing, say that I have direct visual access to what is before me, then the theory works only insofar as it can exploit some justificationist principles. In that case, it would be more accurately described as a reliabilistic justification theory.

### *Reliability and defeat*

Even when the degree of reliability of a belief-grounding process or state seems very high and the process or state is normal, there can be a defeat of would-be knowledge. On this score, the lottery example also challenges reliability theories of knowledge, as it does justificationist theories, and it, too, illustrates the specification problem. Granted, we can characterize the process grounding my belief that I will lose as one in which chance is crucial, and thus claim that the process is not reliable. But as I hold just one out of a million coupons, we might *also* truly describe it as a process that yields true beliefs virtually 100 percent of the time—and we can get as high a percentage as we like by increasing the number of coupons. Under this description, the process sounds very reliable indeed; yet it does not produce knowledge.

Moreover, suppose something like the former description of the belief-grounding process, say, ‘process in which chance is crucial in determining the truth of the belief’, is what reliabilism would use to rule out true beliefs that might otherwise seem to be knowledge. Why should chance not play a role in grounding knowledge? A good answer cannot be that unless we call a belief-grounding process in which chance plays a role unreliable, we cannot

account for knowledge; for that would just assume the reliabilist view that knowledge must be reliably grounded.

In any event, even in perceptual knowledge chance may play a role. It might be by chance that I see you on a passing train: you just happened to be visible to me at the window as the train rushed past. *This* role of chance leaves untouched whatever it is by virtue of which my vision yields knowledge. So how should we specify just what kind of role chance can play in the grounding of knowledge? It may help to say that it can be by chance that we know something but not by chance that, given our grounds for believing *p*, it is only by chance that we are correct. The first is *knowing by chance: fortunate knowledge*. The second is *being right by chance: fortunately true belief*. This distinction is helpful, but still leaves open the question of just what constitutes chance.

There could well be a way around these problems. For one thing, we might say that we often acquire knowledge when the belief constituting it is acquired by chance, as when we just happen to be reminded of something that enables us to finish a crossword puzzle or as in the case of seeing someone on a passing train; but if you know that *p*, it cannot be just a matter of chance that, *given* your believing *p* on the basis you have, it is true, as when we make a lucky guess. Moreover, we might also point out that in the photographic case my belief about Jane's location does not causally *depend* on where she is, since I would believe she is before me even if she were not behind the picture. But this is only the beginning of a solution. For suppose I see her in a mirror, again without knowing that I am not seeing her directly, perhaps because I do not realize that there are trick mirrors at the yard party I am attending. Imagine that she happens to be opposite me, behind the mirror in which I see her, and is reflected into it by other mirrors I do not see (and have no reason to think are there). Here my belief about where she is *would* depend on where she is, since her movements would be reflected in the mirror in which I see her; yet I would still not know that she is opposite me. The kind of causal dependence in question, then, even if necessary for a true belief's knowledge, is not a sufficient condition for that.

The case of my belief that I will lose the lottery is similar. This belief depends on my beliefs about, and in that way may indirectly depend on, the mechanisms that actually result in my losing; but still the belief is not knowledge. The dependence is of course not of the required kind. But now we have another specification problem: how to describe the right kind of dependence, sometimes called a *functional dependence*, but perhaps better called a *discriminative dependence*, since one's belief-forming tendencies are sensitive to fine differences in visual stimuli. If there is a straightforward and illuminating way to specify the right kind of dependence, it is not obvious what it is. A useful metaphor for capturing it, however, is *tracking*. As we track a person in the snow, causally guided by the path, our belief system can be sensitive to the changing evidences that indicate the truth.<sup>14</sup>

### *Reliability, relevant alternatives, and luck*

Even when the appropriate dependence is present just as it normally is, our would-be knowledge can be defeated. To see this, we can alter the deceiving portrait case so that my justified true belief that Jane is before me does have the normal kind of dependence on her location, yet I still do not know she is before me. Imagine that I *do* directly see her standing twelve feet before me and find her looking just as she always does, but this time her identical twin, of whose existence I had no inkling, is a few feet to her right and walking toward the very spot where I see Jane. If I have not learned to tell them apart and would have taken Jane's sister to be Jane had I been ten seconds later, then I surely do not know that Jane is before me.

One way to see that I do not know it is Jane before me is to say that when one knows something, it cannot be just good luck that one is right in believing it, as it is here. But it is not easy to say what constitutes good luck without resorting to something like the notion that you do not have a true belief by mere good luck when your belief is reliably produced—or, perhaps, *sufficiently* reliably produced, or *undefeatedly* justified. Any of these ways of solving the problem will take us back to problems not yet solved.

Some philosophers have dealt with such cases by arguing that the problem in the identical twin case is the existence of a *relevant alternative* to the situation in which there is something, such as Jane's being straight in front of me, that in fact renders one's belief true, an alternative such that one cannot discriminate between the truth of the proposition in question (here, that Jane is before me) and the alternative situation (her sister's being before me). What makes this non-discriminable alternative relevant in the example at hand is the twin's moving toward my field of vision when I first enter, so that I am about to be deceived. On the relevant alternatives view, usually considered a kind of reliabilism regarding knowledge, genuine knowledge is reliably grounded in roughly the sense that the knower can discriminate any relevant alternative from the situation known to exist.

### *Relevant alternatives and epistemological contextualism*

The cases in which the presence of relevant alternatives seems crucial for the question whether someone has knowledge may be conceived in a way that challenges a common assumption. One might naturally assume that there is a single concept of knowledge to be explored by reflection and that its content dictates whether a given true belief constitutes knowledge regardless of the believer's circumstances. Even on this assumption, however, much of our pertinent reflection typically takes place by considering hypothetical situations in which a person might be thought to know something. But again, at least typically these will be approached by considering what may and may not be *said* about the person in these cases, particularly about whether the belief imagined in them constitutes knowledge.

With these points in mind, recall our case of the twins. In cases like this, context seems crucial. Apparently, in the context in which my friend's identical twin whom I cannot tell from the friend is present, I do not know that it is my friend before me. In the context in which the twin is on another continent, I do. Hypothetical circumstances can also create different contexts. Ordinarily I may be said to know that I locked the door when I have the usual recollection of doing so. But what if someone says that there have been burglaries recently and asks if I *know* I did? Here the imagined problem, implicitly recalled by the emphasis on 'know', creates a context in which the standard for achieving knowledge is arguably higher.

The suggested epistemological idea is roughly this. When a "relevant" alternative to what we believe to be the so is actually present in the context and we cannot discriminate it from the case we believe to obtain, our belief does not constitute knowledge even if it is both true and justified. Suppose, however, that the question whether the person knows concerns a context in which such a relevant alternative is simply envisaged. There may then be no question about the person's exhibiting the reliability necessary for knowledge as ordinarily understood, but there may also be an implication that *greater* reliability is needed to warrant *attributing* knowledge. The view that the truth or the falsity attributions of knowledge—including first-person claims to it such as 'I know who did it'—should be judged by different standards in different contexts is often called *contextualism*.<sup>15</sup> But the reliabilist view of knowledge that takes account of changes in reliability in different contexts is also a kind of contextual position. Both views would have us be cautious about claiming to know. But neither view entails that we never in fact have knowledge in some core sense in the kinds of cases in which we unselfconsciously attribute it to people.

To develop a contextualist view, we must be able to decide what alternatives are relevant; the mere possibility of Jane's having a twin in the indicated situation is apparently not—or we would never know our friends are before us (a skeptical view we shall consider in Chapter 13). But suppose Jane has a twin who is never with her, or is currently abroad, or on the way to the party? And does it matter whether we realize there is such a person? These are difficult questions of a kind that a good reliability theory should adequately answer.<sup>16</sup>

Now it may seem natural to conclude that the best way to answer these questions is by considering whether we may properly attribute knowledge in the relevant cases, in which we are now asking about hypothetical uses of language. In this way, philosophy of language seems highly relevant to epistemology. Contextualism is (at least in good part) a position in philosophy of language because it is mainly a view about spoken or otherwise linguistic attributions of knowledge rather than about what constitutes knowledge; but insofar as an account of correct attributions of knowledge bears on what knowledge *is*, the view is epistemologically important.

It should be stressed that contextualism is not the view that there are

different meanings of 'know' in different contexts of ascription, nor the view that such ascriptions are commonly ambiguous (roughly, having two or more meanings and such that one cannot tell which to ascribe). Contextualists stress that the *truth conditions* for ascriptions of knowledge vary with important differences in the context of ascription. They do not generally hold that the meaning of 'know' varies across these contexts. This difference is illustrated by 'now', which varies in what it designates on different occasions of its use ('now' is, e.g., *indexical* in being "indexed" to the time at which it is used), but has the same meaning on each, something like 'the time at the moment of this reference'. Contextualists do not take 'know' to be essentially like 'now' and 'I', but do find differences in the contextual conditions for its correct application. Let us consider kindred views that might be considered contextualist.

That 'know' varies in meaning in different contexts of attribution is, however, implied by a relativistic view on which knowledge-attributions are elliptical, say meaning 'knows relative to everyday standards' versus 'knows relative to high standards'. On this view, one might be said to know, in the first sense, that one will lose in a fair lottery with a million tickets, but not in the second sense, at least when the standards are very high—such as to "rule out error," as even a moderately skeptical person might hold they should.<sup>17</sup> To illustrate, suppose the question is whether I would know that I won't be suddenly able to buy a house as a result of having won a lucrative jackpot. The answer would be positive given 'know' in the everyday sense but negative given 'know' in the "high definition" sense common among skeptics.

On a related view that posits differences in the meaning of knowledge ascriptions, one might take the varying standards for the truth of those attributions to be essentially connected with variations in their meaning. Recall the case of my seeing Jane. If I know her fairly well, then I can know it is Jane I see *as opposed to some other woman of similar age and appearance*, yet not know it is she *as opposed to her twin sister* (who is about to enter the room). On this view—a kind of *contrastivism*—ordinary unqualified knowledge ascriptions are implicitly contrastive and the relevant set of contrasting cases is determinable by contextual features. Thus, to appraise a knowledge ascription we must note not just speaker and context but a contrasting element, such as a proposition other than the one claimed to be known. This is a kind of contextual view, then; but (on the interpretation given here) it accounts for when we may and may not attribute knowledge in terms of differing *meanings* of the attributions rather than just differences in the *standards* for their true application. For a contrastivist, standards *also* vary with differences in meaning; for contextualists, who do not posit different meanings in those attributions, it is just the standards for true attribution that change.<sup>18</sup>

If the relevant alternatives view is a natural route to contextualism, it should not be concluded that the opposing view—*invariantism*—cannot deal with the kinds of cases that challenge any account of knowledge which takes seriously our varying inclinations to attribute knowledge in varying

circumstances. For invariantist views, there is just one set of truth conditions for the statement that one knows that *p*, say knows that one will lose a lottery. The conditions will of course differ from person to person when different people have different kinds and amounts of evidence, but whether one knows will not also vary with context of ascription. Invariantists may argue that “pragmatic” factors, such as the importance of *acting* on the proposition in question, affect the appropriateness of *ascribing* knowledge. They also grant some vagueness in the concept of knowledge, which would account for indeterminacy in whether it applies in a given case, say, when the proposition is that my single ticket will lose a billion-ticket fair lottery.

I do not present any of these problems facing reliability theories of knowledge as insurmountable or any of the contextualist or invariantist solutions as clearly adequate. But reliability theories do face serious difficulties, as do the other theories we have considered. One conclusion that might be drawn here is that knowledge is simply unanalyzable. But that should not be inferred from the difficulties I have brought out. They may be resolvable; and I have of course not discussed all the promising lines of analysis of knowledge there are.<sup>19</sup> It should help us in understanding knowledge to consider other aspects of the concept, and the next chapter will address several important ones.

## Notes

- 1 The most important passages are probably those in Plato's *Theaetetus*, 201c–210b.
- 2 The defeat of this account, which set a spate of detailed critiques in motion, was due to Edmund L. Gettier's now famous ‘Is Justified True Belief Knowledge?’, *Analysis* 23 (1963), 231–3.
- 3 Indeed, Timothy Williamson and others have held that knowledge is the “norm of assertion,” in roughly the sense that one tends to be criticizable for asserting something when one does not know it. See, e.g., *Knowledge and Its Limits*, esp. pp. 238–55.
- 4 If I had believed that she is *directly* opposite me, in the sense that there is no obstacle between us, my belief would have been false; but I would not normally believe that here, where I have no reason even to imagine obstacles.
- 5 The idea that we can learn much, and perhaps the most important things, about the nature of knowledge by determining what renders a true belief knowledge is also used by Alvin Plantinga. Instead of appealing to undefeated justification, he uses a different concept—warrant—which he has characterized functionally as (in rough terms) whatever it is that renders true belief knowledge. For his detailed analysis see esp. *Warrant and Proper Function* (Oxford: Oxford University Press, 1993), particularly chapters 1–3. Extensive critical discussion of this book, and his replies to it, are found in Jonathan L. Kvanvig (ed.), *Warrant in Contemporary*



- Epistemology: Essays in Honor of Plantinga's Epistemology* (Lanham, MD: Rowman and Littlefield, 1996).
- 6 Detailed discussion of how a belief's being knowledge is incompatible with its being only luckily true is provided by Duncan Pritchard in *Epistemic Luck* (Oxford: Clarendon Press, 2005).
  - 7 This of course assumes that we rule out a probability of 1, which is commonly reserved for necessary truths. For a brief and plausible statement of the view that only beliefs based on conclusive (truth-guaranteeing) grounds constitute knowledge, see Fred Dretske, 'Conclusive Reasons', *Australasian Journal of Philosophy* 49 (1971), 1–22.
  - 8 For a valuable treatment of certainty particularly relevant to the one given here see G.E. Moore, 'Certainty', in his *Philosophical Papers* (London: George Allen & Unwin, 1959). Cf. Peter Klein's *Certainty* (Minneapolis: University of Minnesota Press, 1981). I should note that, more than most writers on certainty in epistemology, I am working with the notions central in everyday discourse and not with, e.g., the notions associated with Cartesian epistemology.
  - 9 The thermometer analogy, and some of the impetus for naturalism in recent epistemology, comes from D.M. Armstrong. See esp. his *Belief, Truth and Knowledge* (Cambridge: Cambridge University Press, 1973), though some of the relevant ideas are suggested in his (and others') earlier work.
  - 10 The idea that abstract objects, such as concepts and propositions as standardly construed, are altogether causally inert should not be taken uncritically. For discussion of this issue and a case that no plausible causal requirement undermines the possibility of a priori knowledge, see Plantinga, *Warrant and Proper Function*, pp. 113–17.
  - 11 It seems better to call it reliably *grounded* than reliably *produced* (as many call it), as I could have a belief that is not knowledge until I get appropriate evidence (such as testimony from someone who *is* reliable), and here the belief is not produced by what makes it knowledge, but rather becomes grounded therein. Some writers speak of knowledge as true belief that is reliably produced or *sustained* to capture what I am calling reliable grounding. A further advantage of the wider terminology I use is that it covers both *process reliabilism*, which takes the reliability of process of belief production or sustenance to be crucial for knowledge, and *indicator reliabilism*, which simply requires that the belief be suitably connected with something which indicates that it is true.
  - 12 Plantinga's approach would suggest that the crucial factor in deciding what description to use is how the faculty in question, say vision, is *designed* (whether by God or by evolution or in some other way) to function, for example in judging the identity of persons directly or despite the presence of impervious objects. As our visual faculties are not designed to judge in the latter way, the relevant process is not one



- in which believing someone is in a given place results from seeing the person in a photograph that accurately indicates her whereabouts. One difficulty here is how to determine the design of a faculty. For a detailed presentation of this view see Plantinga, *Warrant and Proper Function*.
- 13 For an approach to solving this specification problem, see William P. Alston, 'How to Think about Reliability', *Philosophical Topics* 23, 1 (1995), 1–29. This paper is in part a response to Richard Feldman's presentation of the problem—under the description 'The Generality Problem'—in 'Reliability and Justification', *The Monist* 68 (1985), 159–74.
  - 14 The metaphor of tracking is usually attributed to Robert Nozick, *Philosophical Explanations* (Cambridge, MA: Harvard University Press, 1981). The most relevant passages are reprinted in Huemer, *Epistemology*, pp. 475–90.
  - 15 For an early discussion of contextualism see, e.g., Keith DeRose, 'Contextualism and Knowledge Attributions', reprinted in Huemer, *Epistemology*. He considers two cases in which he is asked if the bank is open. In one, the matter is routine, in the other, of great importance. He knows it is open in the first, but in the second he does not know without getting further evidence. For "Attributor factors set a certain standard the putative subject of knowledge must live up to in order to make the knowledge attribution true" (p. 497). This is why he can know in one case—in which the context does not demand a high standard for truly attributing knowledge to someone—and not in the other. The implications of this position for skepticism, as well as for understanding knowledge in ordinary cases (on almost any analysis of knowledge we might plausibly give), are significant. Cf. his 'Solving the Skeptical Problem', *Philosophical Review* 104 (1995), 1–52; David Lewis, 'Elusive Knowledge', *Australasian Journal of Philosophy* 74 (1996), 549–67; and (for a critique of Lewis's paper), Stewart Cohen, 'Contextualist Solutions to Epistemological Problems: Scepticism, Gettier, and the Lottery', *Australasian Journal of Philosophy* 76 (1998), 289–306; Jason Stanley, *Knowledge and Practical Interest* (Oxford: Oxford University Press, 2005); and Keith DeRose, *The Case for Contextualism* (Oxford: Oxford University Press, 2009).
  - 16 For a detailed and plausible approach to such questions see Alvin I. Goldman, 'Discrimination and Perceptual Knowledge', *Journal of Philosophy* 73 (1976), 269–82.
  - 17 For detailed discussion of various kinds of relativism concerning knowledge, see Stanley, *Knowledge and Practical Interest*, esp. pp. 138–41 and 144–53.
  - 18 Critical discussion of contrastivism as represented by Jonathan Schaffer is provided by Keith DeRose in *The Case for Contextualism*, pp. 38–41.
  - 19 On Plantinga's notable theory referred to earlier, knowledge is roughly a true belief that results from our cognitive faculties functioning properly

in the relevant situation, which is a matter of their functioning as they were designed to function in such a situation ('environment' may be preferable here to 'situation'). This idea can be adapted to a naturalistic view by construing proper function in biological and psychological terms drawn from a theory of human design, but it is also consonant with a theistic view of our design as determined by divine plan. Plantinga's approach provides a different way of dealing with a number of the problems posed in this chapter (though they remain significant problems), but there is not space to discuss it or other recent theories separately.



# II Knowledge, justification, and truth

## Internalism, externalism, and intellectual virtue

- **Knowledge and justification**

*The apparent possibility of clairvoyant knowledge*

*Natural knowledge*

- **Internalism and externalism in epistemology**

*Some varieties of internalism and externalism*

*The overall contrast between internalism and externalism*

- **Internalist and externalist versions of virtue epistemology**

*Some apparent problems for virtue epistemology*

*The internality of justification and the externality of knowledge*

- **Justification, knowledge, and truth**

- **The value problem**

*Why is knowledge preferable to merely true belief?*

*The value of knowledge compared with that of justified true belief*

- **Theories of truth**

*The correspondence theory of truth*

*Minimalist and redundancy accounts of truth*

*The coherence theory of truth*

*The pragmatic theory of truth*

- **Concluding proposals**

## II Knowledge, justification, and truth

### Internalism, externalism, and intellectual virtue

Whatever the difficulties we face in seeking an analysis of knowledge, our *understanding* of knowledge may be greatly enhanced by considering the nature of its grounds and its relation to other important notions. Indeed, justification is of great epistemological interest even apart from its relation to knowledge. Let us consider justification in ways we have not so far and, in that light, pursue its relation to knowledge.

#### Knowledge and justification

So far, I have often spoken as if, at least in non-testimonial cases, although not all justified true belief constitutes knowledge, all knowledge is at least justified true belief. But if the reliability view is correct in any of its plausible forms, certain true beliefs should be reliably groundable without the subject's having grounds of a kind that yields justification. Let us explore some cases in which knowledge without justification seems possible.

#### *The apparent possibility of clairvoyant knowledge*

Imagine a man who foretells the results of horse races. He always gets them right, even though he never inspects the horses' records, but merely looks at them and their jockeys closely as they amble about and line up. He has no idea why he believes what he does about the results; and after the races he does not even check his accuracy. He does not bet nor especially care who wins. He does, however, have definite beliefs, and we can suppose that it seems natural to him to believe what he does and that there is nothing in his state of mind that would lead him to mistrust himself or think he is guilty of wishful thinking or "crazy." He might be puzzled at forming such beliefs as he does, but he might not even think of their origin or status. It is not clear just how such a thing is possible; but it clearly is possible. There could be a way, for instance, in which both his belief that a horse will win and its actually winning are common effects of the same causes, so that his getting the right answers is not lucky accident, but prophetic in a way, or perhaps sixth-sensory.<sup>1</sup>

Now it appears that this man knows who will win the races. But he surely does not have justified beliefs as to who will win. He *would* have justification if he kept track of his record and noted how well his forecasts turn out. But he does not bother to check on his predictions regularly and has no idea that he is constantly getting the results right. Perhaps we may also assume that (as may be thought to be essential in the case) he also has no good justification for thinking he is not reliable or not justified.<sup>2</sup>

One might protest that he has a kind of foresight which generates directly justified beliefs on the basis of certain experiences, somewhat as perception does. But why would one say this, other than to preserve the view that knowledge implies justified belief? There is no candidate for a sense organ, nor any kind of conscious state that sustains or has any grounding connection with the belief. Nor need he have sensory images representing the victorious horse crossing the finish line; and although we assume that there is some causal process by which he receives the crucial information, we have no idea what it is and cannot plausibly regard it as conferring justification, particularly if the man is puzzled by his having the predictive beliefs at all and has no good reason to think they are justified.

### *Natural knowledge*

Another kind of case (and a more realistic one) argues for the same point. In some of the literature of psychology we read of the *idiot savant*. Such people are considered mentally deficient, yet they have, by nature, as it seems, some extraordinary abilities. We may assume that those adept at arithmetic understand enough to count and to use elementary mathematical concepts. Some of them can apparently just reel off the answers to arithmetical problems that normally require calculation in writing. Let us assume that they regularly get right answers, yet have no idea how they do so: it is not, for instance, by rapidly doing in their head what we would laboriously do in our heads if our memories enabled us to solve the problem mentally. Nor is it by rational insight into the truths in question, such as one has for certain simple logical or arithmetic propositions. It is not known how they do it, and let us assume they have no sense of why they believe the answers in question.

Now consider the first time one of these people—Pip, let us say—reels off the answer to a multiplication problem involving two three-digit numbers. He cannot see that he has a built-in ability or note a series of successes. (We may make a similar assumption about the horse race predictor's initial success.) But he believes the answer and might also know it. For one thing, the belief is a manifestation of an arithmetic ability that is stable and reliable. One could claim there is a mathematical sense that yields directly justified beliefs. But this seems ad hoc, a move designed only to save the view the example counters: that knowing entails justifiedly believing.

If we all turned out to have this mathematical ability under certain conditions, such as an impression of the proposition as true (if only in the form of

an image of it written in boldface), then we might come to believe that there is an arithmetic sense which generates such directly justified beliefs. Perhaps that shows that our concept of justification might evolve; but it does not show that the arithmetic beliefs now in question are justified.

If, as seems likely, Pip's beliefs and those of the horse race predictor are knowledge, they are special cases. We might call them *natural knowledge*, since they seem rooted in the nature of its possessors and do not depend on their training, or their having learned much beyond what is needed to possess the concepts required for holding the relevant arithmetic beliefs, or on their using either their senses or, so far as we can tell, their powers of reason. But even if natural knowledge is rare, its possibility would show that justified belief is not necessarily a constituent in knowledge.

If there can be natural knowledge, that possibility shows something important about both knowledge and justification. What inclines us to grant that Pip knows the answer is chiefly the regularity of his correct results and apparent stability of the mechanism yielding them. The accuracy of the results is surely not accidental; it is rooted in some inner calculative process which regularly—and reliably—yields correct results. On being presented with the problem, he registers the truth. He need have no sense of calculating or even an appearance of truth or self-evidence, such as one may have when one directly grasps an a priori proposition.<sup>3</sup>

There is, then, no mental process of calculation of which the person is aware, or anything else of the sort that grounds justification, as visual impressions can ground it even when one is (unknowingly) having a vivid hallucination. The calculator cannot point to anything to justify the answers, even in the elementary way we can cite how things look to us to justify believing there is a green field before us. This contrast suggests that there may be a major difference between knowledge and justification that explains why the former seems possible without the latter. Let us explore this.

## Internalism and externalism in epistemology

Could it be that justification and knowledge are grounded in quite different ways? Perhaps their grounding differs in a way connected with the basic contrast between them in relation to truth. Apart from self-knowledge, knowledge is at least true belief about the external world (or external matters, such as those of logic). Insofar as it is true belief about the external world, one might expect its grounds to be essentially in that world. The justifiedness of a belief, by contrast, does not entail its truth and seems to many philosophers to rest on a source “inside” the mind.

Some of our examples suggest that what justifies a belief—the ground of its justification—is something internal to the subject. The internal, in the relevant sense, is what we might call the (internally) accessible: that to which one has access by introspection or reflection. Introspection can be simply focusing on what is in consciousness and reflection can be as brief as considering a

proposition. The accessible includes what is actually in consciousness—such as thoughts and visual and other sensory impressions—though here it can be misleading to say that the subject has access, since the phenomena are present to the mind and thus, as it were, being accessed.

The accessible also includes dispositional mental phenomena, such as beliefs and desires. To have (internal) access to something is either to have it in consciousness or to be able, through self-consciousness or at least by reflection, whether introspective or directed “outward” toward an abstract subject matter, to become aware of it, in the (phenomenal) sense that it is in one’s consciousness.<sup>4</sup>

Call the view that justification is grounded in accessible elements *internalism about justification*. By contrast, some of the same examples, such as those of the predictor and the calculator, suggest that knowledge may be grounded entirely in what is external to the mind, in a sense of ‘external’ implying that the grounds are not internally accessible to the subject. Call the view that the grounds of knowledge are at least in part external to the mind *externalism about knowledge*.

There are counterparts of these views: externalism may be maintained for justification, and a kind of internalism may be held to apply to knowledge. The plausible counterparts are not pure (or unrestricted) externalism about justification and pure (or unrestricted) internalism about knowledge. For one thing, because knowledge entails truth it cannot be understood entirely in terms of internal variables, since no combination of these, however well it may justify a belief about the external world, entails the truth of that belief. And there appears to be some respect in which justification is internally grounded, even if it must also imply (say) some objective probability that justified beliefs are true.<sup>5</sup> I propose to say, then, that internalism about knowledge is the restricted internalist view that knowledge is at least in part grounded in elements internal to the mind; and externalism about justification is the restricted externalist view that justification is at least in part grounded in elements external to the mind.

Consider beliefs based on sense experience as plausible support for internalism about justification. My justification for believing that there is a green field before me is grounded in my sense impressions, and I can become directly aware of them (hence access them) by simply “looking within” or by appropriate reflection on my experience. By contrast, the grounds of my knowledge that there is such a field are perceptual and require its actually being there. To become aware that I am perceiving it, by contrast with hallucinating it, I must do more than consider the contents of my mind. I need evidence from other visual standpoints, or from other senses. Even when I get it, skeptics will tend to deny that I know, since I could be repeatedly hallucinating. This skeptical problem will be pursued in Chapters 13 and 14. The point here is simply that sensory experience, as a basis of justification, is accessible to introspection in a way that perceptual experience, as a basis of knowledge, is not.



Justificationist views of knowledge (roughly those that construe it as essentially embodying justification of a kind that is not analyzable in terms of reliability) typically embody an internalist conception of justification. Reliability views of knowledge typically embody a strongly externalist conception of knowledge; and if they appeal to the notion of justification, they conceive it too as grounded in ways that need not be accessible to consciousness, above all to introspection or reflection. (Reflection need not be introspection and is important for the internalist account of our access to the grounds of a priori justification: above all to our understanding of concepts and their relations.)

### *Some varieties of internalism and externalism*

Internalism about justification and externalism about knowledge are, in their qualified forms, compatible, whereas pure internalist and externalist views cannot both hold, either for justification or for knowledge. There are many versions of internalism and externalism, whether they are restricted as I have suggested or meant to apply unrestrictedly to both justification and knowledge.

An important respect in which internalist views differ among themselves concerns how *readily* the justifiers are accessible to consciousness. An important and parallel way in which externalist views differ among themselves is in the kind of non-introspective knowledge or justified belief they take to be possible regarding the grounds of knowledge: one might, for instance, think that commonsense observation is enough to ascertain how reliable perception is, or one might take scientific evidence to be necessary for determining this.

Many points underlie the contrast between internalism and externalism. My concern is chiefly with what seem the most plausible internalist and externalist views: internalism about justification and externalism about knowledge. To simplify matters, let us consider these views mainly in relation to the nature of the grounds of justification and knowledge, not as applied to either *how*, or how strongly, those grounds justify. This is, however, a further respect in which internalist and externalist views differ among themselves as well as from one another, and I will note some points about it.

The imagined internalist about justification holds only that the grounds of one's justified beliefs are internal, for instance sensory states of the kind present in perception or beliefs, of which we can be conscious by virtue of their manifestations in consciousness, such as an assenting thought of the believed propositions. The view does not say that *how*, or how strongly, those grounds justify beliefs based on them must (say, by guaranteeing their truth) be an internal matter and thereby, in principle, accessible to introspection. Similarly, the imagined externalist holds that what grounds knowledge—reliable production or sustenance of the constituent true belief—is not wholly internal, and so not altogether accessible to consciousness, even if part of the

ground, say sensory experience, is. If *what* grounds knowledge is not wholly internal, then *how* it does so is not either.

It is of course natural to think (as reliabilists tend to) that how such belief production grounds knowledge is *less* likely to be accessible to consciousness than what elements ground knowledge. The former is, for one thing, more complex. Similarly, internalists may hold (plausibly) that our access by introspection or reflection to what grounds our justification is better than our access to how it grounds that justification.

### *The overall contrast between internalism and externalism*

If these internalist views about justification and externalist views about knowledge are roughly correct, then the main point of contrast between knowledge and justification is this. Apart from self-knowledge, whose object is in some sense mental and thus in some way internal, what one *knows* is known on the basis of one's meeting conditions that are not (at least not entirely) internally accessible, as states or processes in one's consciousness are. By contrast, what we *justifiedly believe*, or are simply justified in believing, *is* determined by mental states and processes to which we have internal (introspective or reflectional) access: our visual experiences, for instance, or our memory impressions, or our reasoning processes, or our beliefs of supporting propositions. All of these are paradigms of the sorts of things about which we can have much introspective knowledge.

It is significant that for the externalist about knowledge, even introspective knowledge, whose object is mental, is based partly on something not accessible to consciousness, namely the appropriate kind of functional relation between the thing known, say my imaging, and the beliefs about it that constitute self-knowledge, in this case my believing that I am imaging. Roughly, because my imaging process reliably grounds my believing that I am imaging, I know I am; but I have no internal access (and ordinarily none at all) to the reliability of this process. Even if I can be aware of some of the presumably causal connections between imaging and believing I am imaging, I would apparently need inductive, partly external evidence to become aware of the reliability of the process grounding such belief.

What is central for knowledge, in the externalist view, is that the beliefs constituting it register truth. As this objective connection between the grounds of a belief and its truth is understood, the belief-grounding factors are not internal in the crucial way: they are not necessarily accessible to consciousness. An externalist can grant that I can become directly conscious of my believing that I pruned my blue spruce. But internal access to the memorial belief does not imply such access to the reliability of the process by which memory produces true beliefs about the past. I can become aware of that reliability only through a study of how well memory works in producing true beliefs. This requires at least making observations, some of which, like observations of the spruce, are external, and indeed relying on memory of

my results in checking on my memory. I cannot rely on introspection or reflection alone in determining the reliability of my memory; and indeed, I cannot without circularity test it for reliability at all (we will reconsider this last problem in discussing skepticism).

On the other hand, what is central for internalism about justification is that justified beliefs be those that one is in some sense in the right in holding, given the sensory impressions, rational intuitions, and other internal materials introspectively accessible to one. In very broad terms, the strongest contrast may be this. Insofar as we may appropriately use the language of rights to frame an analogy, we might say that the internalist regarding justification tends to conceive justification, in accordance with certain justificational standards, as a matter of having a *right to believe*, and of knowledge as occurring when justification is combined with truth in a certain way; the externalist about knowledge tends to conceive knowledge, in accordance with certain epistemic standards, as a matter of *being right* and of justification as occurring when one's belief is, in a certain way, likely to be right.<sup>6</sup> Just as we can have a right to do something and not do it, we can have justification for believing something and not believe it; but in both cases we have a basis that puts us beyond a certain kind of criticism. In neither case, however, does our right guarantee that we believe truly or do the right thing: in the realm of belief as of action, protection from criticism is not protection from error.

This terminology can be misleading if one thinks of having a right as always applicable to actions; for beliefs are not actions, nor can we in general—if ever—bring it about at will that we believe something, the way we can move our limbs at will. But there are rights to property, and that is not action either. We must also resist thinking that we have a right to believe something when it is not our duty not to believe it—the notion of a right appealed to here is associated with critical standards but does not presuppose ascribing to us duties to believe or withhold belief. There may be some such duties, but even apart from that we should see the notion of a right as facilitating a useful analogy between epistemology and ethics.

The terminology of rights is misleading if one thinks that being within one's rights puts one beyond all criticism. It does not: one might have a right to punish a child but because of special circumstances be criticizable for doing it. Similarly, one could be criticizable for holding a justified belief, say because there is still room for doubt and the matter is so important that one should withhold belief until one's grounds are conclusive. But when one holds a justified belief, there are certain *kinds* of criticisms one is not liable to, such as that of making a groundless assumption or being intellectually sloppy. In different normative language, we might say that justified beliefs are in a certain sense creditworthy, as they are grounded in considerations that support them in a way an intellectually responsible person should recognize. Moreover, just as some rights are better grounded or stronger than others, or both, creditworthiness also admits of degrees.

The central internalist idea about justification is that of meeting a certain

justificational standard that one can conform to on the basis of a kind of response to accessible elements. Internalists strongly associate having justification for belief and a readiness to *justify* it: roughly, to give one's ground(s) for it. This readiness presupposes that the grounds are accessible. This conception of justification contrasts with the reliabilist notion of justification as having, on the right kind of ground, a true or at least probably true belief. The difference might be described as roughly between a kind of permissible, or minimally creditable, belief and a kind of successful belief. Moving further from the language associated with rights (as I think wise), we might say that the contrast is between belief that is internally and justificational well-grounded and belief that is externally (in an objective way) truth-conducively well-grounded. One appeal of the externalist conception of justification is that it links grounds of justification closely with grounds of knowledge, which in some sense seem to be clearly truth-conducive.

### **Internalist and externalist versions of virtue epistemology**

Internalist and externalist approaches in epistemology represent a basic division. The contrast between them can help in understanding any comprehensive epistemological theory, and applying them to a sample theory can also help in understanding them. Consider, for instance, *virtue epistemology*, which represents theories committed to (roughly) the position that knowledge and justified belief are to be understood as expressions of epistemic virtue, taken as a capacity, such as observational acuity, apt for arriving at truth.

Virtue epistemology is in part modeled on virtue ethics, which takes the concept of moral virtue to be the basic moral concept and construes morally sound actions as the kinds that express that capacity, say by being grounded in the virtuous character traits of honesty or justice. For instance, Aristotle said of the virtues of justice and self-control, "actions are called just or temperate when they are the sort that a just or temperate person would do. But the just or temperate person is not [defined as] the one who [merely] does these actions."<sup>7</sup>

Different theories analyze and divide epistemic virtue in different ways, say into observational and a priori virtues and further into perceptual versus introspective virtues and (on the a priori side) into logical and mathematical ones. For an internalist virtue theory, justified belief would be (roughly) belief based on internally accessible grounds understood in terms of, and connected in the person with, an epistemic virtue. Thus, a justified belief might be based on sensory experiences taken as the sort of thing an epistemically responsible person relies on for the kind of proposition in question. On an externalist virtue theory, justified belief would be belief based on processes that are appropriately connected with a virtue and reliably lead to truth. Accurate observations producing true perceptual beliefs would

exemplify such a process, and making them is a typical manifestation of epistemic virtue.<sup>8</sup>

For either kind of virtue theory, knowledge would imply truth; but whereas the internalist would also require that knowledge rest on accessible grounds, the externalist would not. To see the difference better, recall Macbeth's hallucination of the dagger. For an internalist virtue theorist, if his sensory experience is normal enough and he has no accessible reason to doubt his acuity (say because he sees the dagger as hanging on a peg on the wall, rather than—incredibly—as in midair), then his belief that there is a dagger before him may be justified: it is the kind of belief a person with virtuous intellectual character would form in the circumstances. For a reliabilist (hence externalist) virtue theorist, the relevant process grounding his belief is presumably not reliable and his belief does not express epistemic virtue.<sup>9</sup> Similarly, the idiot savant—our lightning calculator—would lack epistemic virtue on the internalist account but might, on some externalist accounts, have it—in the form of a specific arithmetic virtue.

The chief difference is that virtue is defined in terms of internal standards in the first case and external ones in the second. But both views are virtue approaches to justification because they construe it as an expression of epistemic virtue, as opposed to defining epistemic virtue (as many epistemologists would) as the sort of character feature that tends to produce justified belief. It is possible, however, to build internalist requirements into a reliabilist virtue epistemology, just as one may build justification into knowledge. One might hold, for instance, that a belief constitutes knowledge only if it is both reliably produced and manifests an epistemic virtue in such a way that the knower can, with careful reflection (hence by accessing some accessible justificatory element), say something in favor of it.<sup>10</sup>

This contrast between externalist and internalist virtue theories can be developed in many ways. For instance, on externalist lines, the crucial feature of the relevant epistemic virtue would be producing a favorable ratio of true to false beliefs; on internalist lines, the crucial feature would be either producing such a ratio on the basis of internally accessible grounds or—if justification rather than knowledge is the epistemic target—producing a suitable ratio of beliefs that are internally justified.

The contrast between internalist and externalist virtue epistemologies is also expressible in terms of the difference between an *epistemic power* and an *intellectual virtue* in the ordinary sense illustrated by clarity of mind, imagination, and logical rigor. Epistemic powers are exhibited by the lightning calculator who reliably gets correct answers and, similarly, by people who, in the automatic way that some can tell the sex of a chicken, simply detect truths about others' thoughts yet have no idea how they do it. The mathematical and psychological powers in question are certainly good things and, in that generic sense, virtues. But suppose we conceive virtues as Aristotle did and as, in the ethical tradition following him, character traits that one develops by intellectual exercise and through which one acquires justification

on the basis of grounds in one's ken. Then it becomes important to determine whether virtue epistemology is focused more on epistemic powers than, as its name suggests, on intellectual virtues. The determination may be quite different for different positions in virtue epistemology.

### *Some apparent problems for virtue epistemology*

At this point, we can discern an apparent problem confronting epistemic virtue theories. Precisely how can we specify the kind of character feature we wish to call an epistemic virtue without *already* having at least a rough account of justified belief and knowledge? One answer would be that, as in the task of constructing any epistemological theory, we need only an intuitive sense of what count as instances of justification and of knowledge; we can then construct a trait-based account of what they are. Perhaps, however, such a sense is required for construction of any account of justification or knowledge. Consideration of natural knowledge still leads to some doubt whether there is a plausible conception of a virtue, as distinct from a power, whose exercise yields such knowledge. If the lightning calculator knows the answers in an automatic way and with no sense of the basis of the beliefs, there is obviously a cognitive competence and it is clearly a good thing to have. Still, to those who think of virtues as praiseworthy in themselves and as exercised for reasons connected with a kind of good the agent understands, it may seem that the arithmetic competence in question is elevated to a virtue by its desirable outputs rather than being an element in character whose manifestations in beliefs render those beliefs creditworthy.

It is true that we may be able to devise a set of conditions for belief rooted in a good epistemic constitution—epistemic character, in one sense of that phrase—such that virtuously formed belief is equivalent to justified belief (or, when true, to knowledge). The question is whether we do not do better to try to understand the relevant character traits, such as perceptual acuity, intellectual insight, and logical rigor, by appeal to non-virtue-theoretic accounts of justification and knowledge, rather than proceed in the other direction, trying to understand justification and knowledge in terms of intellectual character.

For most epistemologists, the natural approach is first to understand justified belief and knowledge in some non-virtue-theoretic way and then to explicate epistemic virtue as the kind of character trait suited to producing them, a kind, indeed, that can be cultivated by internalizing the more basic standards for appraising belief. Roughly, the idea is that an epistemic virtue is to be analyzed as a trait apt for producing knowledge or justified belief; knowledge is not to be analyzed as, say, true belief reliably produced by an exercise of an epistemic virtue, or justified belief as the kind grounded in an epistemic virtue.

There are various approaches to solving this kind of problem. Quite apart from whether they succeed, the notion of epistemic virtue provides

a distinctive perspective on both knowledge and justification. Moreover, knowledge and justified belief could usefully be understood as *equivalent* to notions rooted in virtue theory, even if the best way to *analyze* the former is along the kinds of internalist or externalist lines we have been exploring.

### *The internality of justification and the externality of knowledge*

Regardless of whether we focus on virtue theories or on the more common epistemological positions we have been considering—individual-belief accounts of justification and knowledge as opposed to epistemic-character accounts—the idea that knowledge is externally grounded and justification internally grounded would help to explain why reliability theories are, in the ways I have indicated, as plausible as they are for knowledge, yet less plausible for justification.

It is true that the sources of justification of belief seem generally to be sources of true belief. But must they be? Could not my apparently normal visual experience in hallucinating a green field sometimes justify me in believing there is one quite as strongly as an ordinary seeing of it? It could. Moreover, though I would not know that there is a green field before me, the internalist would hold that my justification for believing there is could be quite as good as it would be if I did know it.

To be sure, if I justifiedly believe I may be hallucinating, then I am unlikely to be justified in believing there is a green field there. But my beliefs, including beliefs about possible hallucinations, are themselves internal. We thus have one internal factor affecting the way another bears on justification, not an external factor preventing a basic source of justification from generating it. Here, then, internalism does justice to the phenomenon of defeat of justification.

Moreover, notice how the clear cases of highly reliable belief production illustrated by the predictor and the lightning calculator do not appear to generate justification, though they do appear to generate knowledge. Furthermore, no matter how reliable my perceptual processes are, say in giving me impressions of birds flying by, and thereby true beliefs that they are flying by, if I confidently and reflectively believe that my vision is unreliable, and especially if I also justifiedly believe this, then it is doubtful that I am justified in believing that birds are flying by. The more confident and reflective my justified belief that my vision is unreliable, the less the justification, if any, of my belief that birds are flying by. Thus, in addition to reliable grounding alone not producing justification, its apparent capacity to produce justification in common circumstances is vulnerable (at least) to justified beliefs that the beliefs it produces are unjustified or their underpinnings unreliable.

If knowledge and justification do contrast in the suggested way, why is justification important to knowledge at all, as it certainly seems to be?



Here are two points by way of partial answer. First, the sources of justified belief—experience and reason—are generally (if in a different way) sources of knowledge. Second, virtually the only knowledge we can conceive of for beings like ourselves—and certainly the normal kind we take ourselves to have—is apparently grounded, at least indirectly, in those sources, broadly understood. If these points are correct, then we can at least understand how knowledge typically arises if we think of it as (in part) justified belief. If, moreover, we think of it as *appropriately* justified true belief, then, conceiving knowledge under that description, we can at least pick out the vast majority of its instances.

## Justification, knowledge, and truth

There may be a further, perhaps deeper, point implicit in what has been said about justification and knowledge. Justification by its very nature has some kind of connection with truth. One can see this by noting that there is something fundamentally wrong with supposing that a belief's being justified has nothing whatever to do with its truth. This in turn can be seen by considering how the *process* of justifying a belief, conceived as showing that the belief has the *property* of being justified, is always taken to provide grounds for considering the belief true. Justification (justifying) of our beliefs is by its nature the sort of thing we do when their truth is challenged; their justifiedness—which entails justification for taking them to be *true*—is what this process of justification shows when it succeeds.

The connection between justification and truth is perhaps most readily seen in the case of a priori justification. Consider such paradigm cases as a priori justified beliefs of self-evident propositions and of propositions that self-evidently follow from them, as do some very simple theorems of logic. Here, our apparently most plausible conception of a priori justification entails the truth of the beliefs so justified.<sup>11</sup> These cases are unlike perceptual ones in that if a belief claimed to be a priori justified turns out to be false, there is at least normally a defect in the purported justification, say a misunderstanding or a careless error in reasoning, whereas a false perceptual belief can be strongly justified.

Justification of empirical beliefs also seems connected with truth. If, for instance, I am justified, by a clear visual impression, in believing there is a field before me, then I may take it to be true that there is one. If, on the other hand, we discovered that a certain kind of empirical belief is always false, we would not consider a belief of that kind justified. Imagine that the smell of onions ceased to indicate their presence and that beliefs grounded in it no longer correspond to the facts (thus to truth) as determined by other sources of belief, such as vision and touch. Then we would have good reason to cease to regard these olfactory impressions as a source of justification. As we cease to take a ground of belief as indicating truth, we tend to cease to take beliefs thus grounded to be justified.



These points about the relation of justification to truth suggest that even if it is an internal matter whether a belief is justified, the standards we use for determining justification are responsive to our considered judgments about which internal sources tend to produce true beliefs. The way we conceive justification, then, makes it well suited to help us understand knowledge, in at least this respect: when a belief is justified, it has the sort of property which, by its very nature as apparently grounding the belief in the real world, we take to count toward the truth of the belief, hence (other things being equal) toward its being knowledge.

Justified true belief need not be knowledge, and knowledge apparently need not be justified belief. But normally knowledge arises from the same sources as justification: normally, the internal states and processes that justify our beliefs also connect our beliefs with the external facts in virtue of which those beliefs are true.

## The value problem

In the light of what we have now seen regarding the nature of knowledge and the relation between knowledge and justification, we are in a position to explore what has come to be called the value problem—one raised by Plato in his dialogue *Meno*. Why is knowledge more valuable than mere true belief? This is a difficult and multi-dimensional question. Only a few points pertinent to this problem can be made here.<sup>12</sup>

### *Why is knowledge preferable to merely true belief?*

It is important to see at the outset what kind of value is at issue. Two kinds are commonly distinguished, the kind of value things may have *in themselves*—commonly called *intrinsic value*—and the kind of value things may have as a *means* to something of intrinsic value. The latter is typically called *instrumental value* (and a thing may have both kinds of value). We might now ask whether knowledge has value in itself, whether it has instrumental value, and whether it has either kind *as* knowledge rather than simply as true belief.

There is some plausibility in saying that knowledge is both good in itself and (often) good as a means to other things that are good in themselves, such as a beautiful building whose creation knowledge makes possible or, to shift from product to process, enjoying the construction of such a building. As these examples show, there is a profound difference between kinds of things that are good in themselves: enjoying something is an experience; but neither buildings nor knowledge is an experience. I propose to call non-experiential things that are good in themselves *inherently good* and treat only experiences as strictly speaking intrinsically good.

One reason to make this distinction is that intrinsic goodness is valuationally more basic than inherent goodness (and is plausibly taken to be the basic kind of value). Consider a beautiful painting. It is good in itself,

non-relationally, though it may *also* be good relationally, say as a means. Beautiful things are good whether or not they are instrumentally good, as means to some further good. Still, the beautiful painting is good *because* of the intrinsic goodness of appropriate (actual or possible) experiences of its colors and shapes. If no one has a good (aesthetic) experience of the painting, its inherent value is in a sense unrealized. The point of painting is (largely) to produce something that yields a good experience in being properly viewed. The painting is not, however, a means to this experience; it is a *constituent* in it, not, say, a cause of it. As this suggests, in such cases we should conceive the inherently good not as good on account of what it is a means to, but as good depending on whether an appropriate kind of experience of it on the basis of its non-relational properties is intrinsically good, for instance aesthetically rewarding.<sup>13</sup>

To ask, then, whether knowledge is good in itself is implicitly to ask about the intrinsic value of an appropriate kind of experience of it, say in comprehending contemplation of knowledge of a logical principle. This is not merely contemplating a belief that *is* knowledge, but contemplating a belief *as* knowledge, in this case as knowledge of the principle in question. With everyday knowledge, such as that there is furniture around me, it is not obvious that contemplating it, even as knowledge, is inherently valuable. But there can be a kind of positive and valuable response to contemplating one's knowing such things. In any case, *some* instances of knowledge surely are inherently valuable; and even apart from that it appears that *if* merely true belief that *p* has any inherent value, then (other things equal) knowledge that *p* has more. This approach is helpful in answering the question why merely true belief might not have the same kinds of value as knowledge and whether, even if not, *justified* true belief might. Let us take these questions in turn.

Mere true belief, belief which is not even knowledge or justified, does not seem inherently good, say creditworthy. Reflect on what it typically is: only by accident or by good fortune is the belief true. The believer is in no way creditworthy for being right, and it is plausible to think that, as manifesting this point (and perhaps in other ways), the belief is not creditworthy either. Moreover, there is some reason to think that—at least from the point of view of being intellectually virtuous—one ought not to hold beliefs that are not (to at least some extent) justified. Reflection on this kind of belief does not have the positive character (a kind of intrinsic goodness) that should mark comprehending reflection on something good in itself.

The verdict regarding the instrumental value of merely true beliefs seems more complicated. On the plausible assumption that true beliefs by and large guide our lives better than false ones, we might say at least this. First, although a true belief—say that eating raw pork is risky—may produce or guide the same actions whether it is justified or not, if it is not based on a justifying ground, it will not, other things equal, be as stable as if it were so based; one would not, for instance, have any basis for it that would lead one to resist giving it up under even slight (misleading) counter-evidence.

Granted, other things may not be equal. If, for instance, I justifiedly believe that  $p$  only on the basis of testimony, you could undercut my justification by plausibly attacking my source even if I retained my belief that  $p$ . But many kinds of grounds for justification (as for knowledge), such as the many kinds provided directly by perception or intuition, are not easily undercut or overridden; and in any case, other things equal, beliefs that are justified or constitute knowledge are more stable than beliefs that are true but do not have a justificatory or an epistemic basis.

To be sure, a bad reason for believing  $p$ —such as a sophistical consideration that does not even probabilistically support  $p$ —might happen to give the belief much stability. But believing for a bad reason would not only be an inherently bad thing, in being a condition contrary to intellectual virtue, which it ill-befits. It would also produce an at least instrumentally bad tendency to believe for bad reasons of the same kind. Believing for bad reasons may, more often than not, be expected to imply or lead to believing falsehoods. Second, if one has no adequate support for believing  $p$ —which would be typical with beliefs held for bad reasons<sup>14</sup>—then, other things equal, one cannot as easily get others to accept  $p$ , and this may interfere with harmonious relations, or at least reduce the chance of getting others' support.

### *The value of knowledge compared with that of justified true belief*

The value which knowledge that  $p$  tends to have over that of *justified* true belief that  $p$  is more difficult to see. But recall the cases in which one has a justified true belief that is not knowledge. In many of them, one is correct only by accident or good fortune. In that case, the belief would seem to have at best less inherent value than knowledge, even if it is still creditworthy to some extent, depending on the kind and degree of justification. But what of the lottery case? There a justified belief that one will lose is not true by accident or good fortune, since one has only a single ticket out of a million. Still, knowing one will lose would require a basis beyond the mere probability calculation underlying one's justification. This basis would yield additional grounding and would tend to provide additional stability. Stability in a true belief seems inherently valuable, at least where we see the stability as based on knowledge-grounding elements. Such stability is surely also instrumentally valuable. True beliefs help us to map reality; they should not be easily erased. Justifiers and knowledge-grounding elements tend to sustain true beliefs based on them.

A further point about at least many justified true beliefs, and hence about much knowledge, is that they are intimately connected with whatever measure of understanding goes with their particular justification. This may not be appreciable, as where testimony is one's only basis of justification; but for much that we justifiedly believe, our justifying ground also provides some understanding—indeed that ground *is* understanding in the case of those

self-evident propositions we justifiedly believe in the non-inferential way described in Chapter 5. Similarly, if I believe  $p$  on the basis of proving it, my proof both justifies my belief and provides a kind of understanding of why  $p$  is true. By contrast, in at least many cases of empirically justified belief, we gain or can gain some understanding of the justifying factor from the proposition it justifies for us. If knowledge of a distinctive roar is my basis for believing a helicopter is flying nearby, I can understand the character of the sounds I hear on the basis of taking them to be caused by engine and propeller sounds. Very roughly, the difference in question is that in the *a priori* cases what justifies provides understanding of the content of the belief it justifies, and in the empirical cases the content of the justified belief provides understanding of the justifier. May we say, however, that other things equal, where one knows  $p$  as opposed to simply having a justified true belief of it, the understanding provided in the ways just illustrated is *greater*? This is far from obvious and may not be so, but it might still be true that in the former case, the understanding is more stable.

I have so far presupposed that understanding has value. It surely does. Understanding, even when dispositional, is inherently as well as instrumentally valuable. Contemplating it tends to have a positive element and may embody the reward of satisfaction. We seek understanding for its own sake—hence non-instrumentally. Contemplating it often embodies the sense of having achieved—or at least of possessing—something worthwhile. The instrumental value of understanding is perhaps even more variable than its inherent value, but understanding anything can be a means to achieving something.

Much of our understanding is dispositional and as such not in consciousness (though we may be conscious *that* we have it). Consciousness that I see why knowledge is not justified true belief is not consciousness *of* something that explains why not. But understanding may also be occurrent, and occurrently understanding a proposition, as when one consciously realizes the full meaning of a theory, seems *intrinsically* good. There is often a special gratification in the occurrent understanding that consists in consciously seeing *why* something is so. “Ah, of course!” one may sometimes think triumphally.

Regarding instrumental value, knowledge again generally adds something to mere justified true belief. Take the lottery. If I know I will lose, this must be on a ground, say because I know the drawing is fixed. Here (and in similar cases) I would have more to say to others about why I believe I will lose than where I believe this from calculating the probabilities. This would add instrumental value to my belief, for instance where I need to eliminate someone’s fervent hope of my soon becoming rich by winning the jackpot. (Some hope would be sustainable in a fair lottery.) I would also have a better explanation of why I believe I will lose than if I am justified simply by probability calculations. I have an additional explanatory factor and can now account for how my explanation of why I will lose differs from one that, before the drawing, the winner could give.

An even more difficult challenge occurs if we countenance knowledge that is constituted by *non*-justified belief, as with the lightning calculator. Perhaps this kind of knowledge is not better in itself than justified true belief. But it may be more accurate to say that justified true belief is better than such knowledge on the basis of manifesting intellectual virtue, but not as good on the basis of reflecting or otherwise reliably indicating the *facts*. The arithmetic beliefs in question, for instance, are based on the deliverance of a reliable mechanism. In this case, however, the calculator might not know why *p* is true, whereas those who calculate the result in the normal way can explain, in terms of arithmetic principles, why it is true.

As to instrumental value, we again find mixed results. Which of the two, mere justified true belief and knowledge constituted by non-justified belief, is more stable would vary from case to case, but even if the non-justified beliefs constituting knowledge are deeply rooted and highly stable, they might be difficult to justify even upon inquiring into possible evidence for them, as with the horse race predictor. They would then lack the kind of instrumental value belonging to beliefs of which we can convince others by citing our justification, or for which we can at least provide them—and ourselves—with the rational basis that our justification constitutes.

There are, then, some good reasons for several tentative conclusions. (1) At least where knowledge embodies justified true belief, it is inherently better than mere true belief, and (2) justified true belief that does not constitute knowledge is also inherently better than mere true belief. (3) Both points are supported by the role of justification in yielding understanding, which is an inherent good. (4) By and large, the same preferability expressed in (1) and (2) holds for instrumental value, and (5) where knowledge does not embody justification, it may be inherently good on some counts and inherently bad on others (e.g. intellectual virtue). But (6) such knowledge is likely to be instrumentally better than mere true belief, though (7) perhaps not instrumentally better than *justified* true belief that does not constitute knowledge.

## Theories of truth

The notion of truth figures in much of this book, and it has been discussed voluminously by philosophers. But even though it is more a metaphysical problem than an epistemological one, an outline of some of the main positions on the nature of truth may help us. It will add definiteness to the major theories of knowledge we have considered, and, more important, it will enable us to see that they can be largely neutral on the difficult question of precisely what account of truth is sound.

### *The correspondence theory of truth*

The way of speaking of the truth that seems most natural in epistemology suggests that truth, like knowledge, is external (apart from propositions with such internal content as that I am thinking about knowledge). Whether there

is a green field before me is not a matter of states of my mind. It seems to be an objective matter independent of anyone's mind, and the green seems to be present or not regardless of whether we believe it is. Whether my belief is true is determined by whether the field is actually there: by external reality.

Sometimes this is put by saying that in general the truth of our beliefs is not mind-dependent. If truth is not mind-dependent (unless it is about mind, say that I am silently reciting a line of Shakespeare), and if truth is at least in that sense objective, then we have a version of *realism*, roughly the view that there are external things which are as they are independently of what we take them to be. I am thinking of true propositions, and of truth as represented by them, whether believed or not, along the lines of a version of the correspondence theory of truth. Its central thesis is that true propositions "correspond" with reality. It is usually added that they are true in virtue of that correspondence. Thus, the proposition that there is a green field before me is true if in reality there is a green field before me; and it might also be said to be true in virtue of there really being such a field before me.<sup>15</sup>

An expression apparently equivalent to the first, modest formulation of the correspondence view would be this: to say that the proposition is true is to say that it correctly represents reality. This, in turn, is commonly taken to mean that it represents a fact. How else could we even think of truth, one might wonder? What else could it mean to say that a proposition is true other than that things (or the facts) really are as the proposition says?

Agreement on this, however, leaves room for much diversity. One can think of facts metaphysically or, in an ideal sense, methodologically. An important example of the latter view is the idea (articulated by the nineteenth-century philosopher C.S. Peirce) that truth is what scientific investigators would ultimately agree on. On the most common version of the metaphysical view, by contrast, there can be truths that would evade inquiry for ever, and so the concept of truth is not analyzable even in terms of an ideal method.

### *Minimalist and redundancy accounts of truth*

One answer to the question of how to conceive truth uses the schema '*p*' is true if and only if *p*; for instance, 'Grass is green' is true if, and only if, grass is green.<sup>16</sup> One might now argue that to say something is true is equivalent to asserting it, in the way illustrated, and go on to hold that this equivalence is at least the main thing we need to understand about truth. This approach is associated with a minimalist account of truth. The idea is roughly that there is no more to understanding what constitutes truth than understanding this equivalence. For instance, we know what it is for people to assert propositions; we normally know what kinds of considerations confirm or disconfirm those; and we know under what sorts of conditions to agree or disagree.

If the minimalist account is correct, then asserting the correspondence of true propositions either with "reality" or with "the facts" is nothing more than an equivalent of endorsing the schema, or it goes too far. One might think that to give such a minimalist reading of the correspondence view is

to abandon it altogether. Perhaps that is so, but at least the schema expresses a kind of correspondence: an equivalence between calling a proposition true and asserting it. There is a sense, moreover, in which each of these affirmational acts can itself correspond to facts.

One might go even further than the minimalist account and say that, for instance, ‘“Grass is green” is true’ is not just equivalent to ‘Grass is green’ but has essentially the same meaning. The relation would thus be like that between ‘circle’ and ‘locus of points equidistant from a given point’ rather than like that between ‘circle’ and ‘shape whose circumference equals its diameter times pi’. This identity of meaning, unlike the equivalence asserted by the minimalist view, would make the phrase ‘is true’ redundant in predicating it of a proposition. ‘It is true that  $p$ ’ would add nothing to the meaning of simply asserting that  $p$ , though it might be an emphatic way to assert it. On that basis, one might speak of a redundancy account of truth. On this view, saying ‘That is true’ is another way, perhaps a sometimes more emphatic way, of saying that  $p$ . But it has exactly the same content.<sup>17</sup>

### *The coherence theory of truth*

There are other alternatives to the correspondence view. The best known is perhaps the coherence theory of truth. It takes many forms, but its central idea, expressed broadly, is that a true proposition is one that coheres appropriately with certain other propositions. (The theory may also be expressed in terms of what it is for beliefs to be true, but that formulation invites confusion of the coherence theory of truth with the coherence theory of knowledge, which, though knowledge is constituted by belief, is a quite different theory and does not depend on the coherence theory of truth.)

I cannot discuss truth in detail, but let me indicate how a coherence theory of truth might go if justification is its central concept. In outline, it might say that a true proposition is one which is fully justified by virtue of coherence with every other relevant justified proposition, where a justified proposition is, minimally, one that at least someone is (or anyway might be) justified in believing.<sup>18</sup>

There are serious difficulties in determining what propositions are both justified and relevant to the truth of another proposition which is plausibly thought to be true in virtue of coherence with them. Perhaps a plausible example of how truth can be based on coherence would be a proposition I am perceptually justified in believing, say that there is a maple tree before me, which coheres with what I justifiedly believe on the basis of memory, introspection, inference, and so on, as well as with what I or others would be justified in believing in these ways. This proposition would be true in virtue of coherence with others, such as that I seem to remember a maple there. To say that it is false, by contrast, would be to call it incoherent with certain others, such as the proposition that I fail to have a visual impression of a tree in the relevant place.



The propositions for which I now have justification are not the only ones that matter. If they were, then if I visually hallucinated a maple tree systematically enough, say with accompanying tactual hallucinations and supporting memory impressions, it would be true that there is one before me. By making the set of relevant propositions indefinitely large, the theory seeks to prevent such embarrassing results. Thus, if I am hallucinating, there is surely some proposition I could come to be justified in believing, say that the “tree” will not burn—something I might discover by trying to ignite it—which is not coherent with the proposition that there is a maple there.

Suppose, however, that owing to some quirk of nature—or to some evil genius of the kind Descartes imagined in his skeptical passages—there is no proposition I could become justified in believing that is incoherent with there being a tree before me. If, for instance, I take a torch to the “foliage,” I hallucinate flames. If the evil genius undetectably ensures that every such test is positive, the coherence theorist seems forced to conclude that it is true after all that there is a tree before me. But surely it is still possible that I am merely hallucinating and that it is false that there is a tree there, despite the unending series of justified beliefs I have, or can have, confirming that there is one.

This kind of possibility has led critics of the coherence theory to say that the truth of a proposition is simply not exhausted by our coherent beliefs or potential beliefs supporting that proposition, even when they are justified. Another way to put it is to say that truth is not a construct out of evidence, even excellent evidence that produces a coherent body of beliefs. We can better understand this point if we consider a related theory of truth.

### *The pragmatic theory of truth*

There is also a negative motivation for the coherence theory of truth. When we try to understand what correspondence means, we seem thrown back on some kind of coherence. To say that the proposition that the tree is green corresponds with reality seems to come to little more than saying that in testing this proposition, say by examining the tree in good light, one will always get (or will at least in the main get) confirming results, that is (one might argue), discover propositions that cohere well with the original one. For instance, boiling its leaves will produce a green broth.

This kind of point has led some thinkers to go further and hold a pragmatic theory of truth, on which true propositions are simply those that “work,” in the sense that they are successful in practice—pragmatically. What this comes to is chiefly the claim that believing them, acting on them, and otherwise confirming them, leads (at least in the long run) to positive results, such as spectrographic confirmation of the tree’s color.<sup>19</sup>

Certainly we do not expect a genuine truth to fail us. If, for instance, there really is a maple there, then I can find shade under it, cut wood from it, and expect others to verify its presence. And we *do* expect falsehoods to fail us



eventually. With enough testing, we tend to think, there will (in principle) be disconfirmation. What more is there to truth as correspondence or, for that matter, truth as coherence than such pragmatic success—especially if we consider it over a sufficiently “long run”?

Correspondence theorists have replied that points made by proponents of coherence (and pragmatic) theories of truth confuse the *criteria* of truth, roughly, the standards for determining whether a proposition is true, with the *nature* of truth, what it is. Turning blue litmus paper red is a criterion of acidity, but it is not what acidity is; that must be characterized in theoretical terms involving, say, atomic structure. In support of this, they often argue that a false proposition could cohere with all the propositions that are ever justified, including those discovered in attempted confirmation of it. We might, after all, be permanently unlucky in testing it—or permanently foiled by an evil genius (a possibility pursued in Chapter 13)—so that we never discover its falsity.

These points parallel some made against phenomenalism, which may (though it need not) be held by proponents of either a coherence or a pragmatic theory of truth. In general, a sufficiently powerful evil genius might prevent one from discovering that a stable, recurring set of sense-data which coheres with one’s other sense-data derives from hallucination and does not represent a concrete object. If it is possible for coherence to be systematically misleading in this way, then neither coherence with justified propositions nor any other kind of pure coherence can be what truth is.

I cannot pursue this issue, but it should be plain that it is crucial to assessing the pragmatic and coherence theories of truth. I want to add only that despite the similarities between the coherence theory of truth and the coherence theory of justification, neither theory entails the other. The analysis of knowledge, moreover, can be discussed within either framework for conceiving truth.

It appears, however, that particularly if one favors a reliability theory of knowledge, the correspondence view of truth, even on a minimalist interpretation, seems most appropriate. This is in part because the notion of reliable production is at least not readily analyzed along coherentist or pragmatic lines, especially if justification is central in truth as the coherence and pragmatic theories of truth conceive truth. For then the apparently value-laden notion of justification would be required for understanding reliability, which is characteristically conceived by reliabilists in naturalistic terms, in part as a property belonging to processes that produce true beliefs. None of this implies the falsity of coherence or pragmatic theories in every form, but perhaps enough has been said to suggest some presumption in favor of *some* version of the correspondence theory.

## Concluding proposals

Is there no analysis of knowledge that we may tentatively accept as correct and illuminating? There certainly may be; the ones I have discussed are only

a representative sample of the available analyses, and even they can be refined in response to problems of the kinds I have raised. But there may be no simple and straightforward analysis of knowledge which is *both* illuminating and clearly correct. Much depends on how detailed an account must be to count as an analysis.

We may be able to formulate what is at least a sound *conception* of knowledge, and this should help in seeking a full-dress account. We might say that *knowledge is true belief based in the right way on the right kind of ground*. This conception leaves a great deal open, but what we have seen in this chapter and earlier ones indicates many ways in which we might develop the conception into a detailed account.

The conception leaves open that it may, but need not, turn out that the right kind of basis of knowledge is in part causal. It may, but need not, turn out that the right kind of ground always justifies the belief. It may, but need not, turn out that the right kind of ground is always accessible to consciousness, or is a reliable producer of true beliefs, or is expression of an epistemic virtue. And it may, but need not, turn out that ultimately epistemic chains terminate in experience or reason, or in some other kind of ground of knowledge which is direct in the way foundationalism maintains it is.

In a similar vein, we might conceive justified belief as *well-grounded belief*. Like knowledge, justified belief must be based in the right way on the right kind of ground. We have seen what at least some of the appropriate kinds of grounds are: most basically, perceptual, memorial, introspective, and rational, but sufficient grounds may also be testimonial or inferential. However, the conception leaves open the same kinds of things as the conception of knowledge: whether the right kind of ground is in part causal, whether it is accessible to consciousness or simply a reliable source of true beliefs, whether it is (or is suitably connected with) an epistemic virtue, and whether justificational chains terminate in experience or reason, or in some other kind of ground of knowledge, that is direct in the way foundationalism maintains it is. In exploring justification and knowledge, I have exhibited some of these options as preferable to others, but here my point is simply that the suggested conceptions of knowledge and justification provide a good focus for inquiry regardless of our position on these options.

Quite apart from how these broad questions about knowledge and justification are resolved, then, the conceptions just sketched indicate where a great deal of the work in understanding knowledge and justification must be done. We need an account of *how* knowledge and justification are based on whatever it is in virtue of which they count as knowledge or justification, for instance perception, introspection, and reason; and this will require an account of the inferential transmission of knowledge and justification as well as of their non-inferential grounding. Here I have suggested a partly causal account of both inferential and non-inferential grounds and a moderately holistic account of inductive and deductive transmission of justification and knowledge. We need an understanding of whether the appropriate bases of knowledge must ground it through generating justified belief, or may yield

knowledge independently of justification. Here I have suggested that an internalist account seems preferable for justification and a qualifiedly externalist one seems preferable for knowledge. We also need a general understanding of what it is for a belief constituting knowledge to be true. And we need an account of whether the ultimate grounding of knowledge and justification is some kind of coherence among one's beliefs or, as seems more likely on the basis of what has emerged in this book, anchoring in experiential and rational foundations.<sup>20</sup>

## Notes

- 1 This example is a variant of the central case described in D.H. Lawrence's story, 'The Rocking Horse Winner'. If the case presented seems too far from real possibility, the same point can be made using one in which it is easier to imagine relevant causal mechanisms, say when someone foresees rain by the effect of sensations that in effect measure meteorological conditions.
- 2 One might argue that it is obvious that such a belief must be unjustified. I can see a case for its being so, but I do not find it compelling. I also doubt that his having justification for believing his predictive belief unjustified must prevent its constituting knowledge.
- 3 As this description suggests, I am not thinking of the relevant knowledge as a priori even in the loose sense, but perhaps a case can be made that a kind of understanding of the relevant proposition, as opposed to, say, a subliminal calculation, grounds the knowledge.
- 4 These examples of accessible items suggest that internalism might be characterized as the mentalist view that justification is grounded in *mental* elements, and indeed Earl Conee and Richard Feldman have argued for that. See esp. their 'Internalism Defended', in Hilary Kornblith (ed.), *Epistemology: Internalism and Externalism* (Malden, MA: Blackwell, 2001). It may be plausibly argued that what I call internally accessible is mental, at least if we omit abstract entities such as concepts. But in part to include the abstract and in part for other reasons, I prefer not to assimilate the accessible to the mental.
- 5 For two quite different externalist compromises with internalism see Alvin I. Goldman, 'Strong and Weak Justification', *Philosophical Perspectives* 2 (1988) 51–69, and William P. Alston, 'An Internalist Externalism', in his *Epistemic Justification* (Ithaca, NY: Cornell University Press, 1989).
- 6 A.J. Ayer is widely known for having taken knowledge (and I think the kind of justification appropriate for it as well) to entail a "right to be sure." See *The Problem of Knowledge* (Harmondsworth, Middlesex: Penguin Books, 1956), esp. chapter 1. For further discussion see R.M. Chisholm, *Theory of Knowledge*, 2nd edn (Englewood Cliffs, NJ: Prentice-Hall, 1977), and Carl Ginet, *Knowledge, Perception, and Memory* (Dordrecht:

- D. Reidel, 1978), esp. chapter 3. Extensive discussion of the deontological conception of justification is provided in Matthias Steup (ed.), *Knowledge, Truth, and Duty* (Oxford: Oxford University Press, 2001).
- 7 *Nicomachean Ethics* 1105b6–8, trans. by Terence Irwin (Indianapolis: Hackett, 1999). The bracketed ‘defined as’ is my insertion, the bracketed ‘merely’ Irwin’s.
- 8 Virtue epistemology, despite its roots in ancient Greek philosophy, has received attention in epistemology only in recent years. For a sense of the development and variety of the perspective, see esp. Ernest Sosa, ‘Knowledge and Intellectual Virtue’, in his *Knowledge in Perspective* (Cambridge: Cambridge University Press, 1991); Jonathan L. Kvanvig, *The Intellectual Virtues and the Life of the Mind* (Lanham, MD: Rowman and Littlefield, 1992); James Montmarquet, *Epistemic Virtue and Doxastic Responsibility* (Lanham, MD: Rowman and Littlefield, 1993); Linda Zagzebski, *Virtues of the Mind* (Cambridge: Cambridge University Press, 1996); John Greco, *Putting Skeptics in Their Place* (Cambridge: Cambridge University Press, 2000); Abrol Fairweather and Linda Zagzebski (eds.), *Virtue and Duty* (Oxford: Oxford University Press, 2001) (this contains papers by most of the philosophers mentioned as well as my ‘Intellectual Virtue and Justified Belief’, which develops ideas about virtue epistemology that I suggest here), and Ernest Sosa, *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, vol. 1 (Oxford: Oxford University Press, 2007).
- 9 I say ‘presumably’ because it might be that sensory experience *is* a reliable ground, because of how rarely it misleads, say because hallucinations either are uncommon or (because they do not feel normal) do not produce belief.
- 10 We thus find Sosa saying (in *A Virtue Epistemology*), “Knowledge is apt performance in the way of belief” (p. 41), where “A performance is apt if, and only if, it is correct attributively to a competence exercised by the performer, in conditions appropriate for its exercise” (p. 92). I take the relevant kind of competence to be a kind of virtue. See also pp. 22–42.
- 11 This entailment thesis is obvious if we adopt the conceptions of the self-evident and of strict a priori justification proposed in Chapter 4. But simply working with the intuitive notions of the self-evident and the a priori, the entailment claim has some plausibility.
- 12 For wide-ranging discussion of the problem and an extensive bibliography concerning it see Duncan Pritchard, ‘Recent Work on Epistemic Value’, *American Philosophical Quarterly* 44, 2 (2007), 84–110 and Jonathan Kvanvig, *The Value of Knowledge and the Pursuit of Understanding* (Cambridge: Cambridge University Press, 2003). For further discussion see Marian David, ‘Truth as the Primary Epistemic Goal: A Working Hypothesis’, in M. Steup and E. Sosa (eds.), *Contemporary Debates in Philosophy* (Oxford: Blackwell 2005), and Ram Neta and Duncan Pritchard (eds.), *Arguing about Knowledge* (London: Routledge, 2009),

- Part Two, which contains *Meno* and recent papers by John Greco, Jonathan Kvanvig, and Linda Zagzebski.
- 13 I develop and defend this distinction in chapter 4 of *The Good in the Right: A Theory of Intuition and Intrinsic Value* (Princeton, NJ: Princeton University Press, 2004) and further in 'Intrinsic Value and Reasons for Action', in Terry Horgan and Mark Timmons (eds.) *Metaethics after Moore* (Oxford: Oxford University Press, 2006).
  - 14 A person's having no adequate support for *p* would be only *typical* of such cases because one might *have* some good reason to believe it yet simply not have noticed it or come to believe *p* for the reason(s) in question.
  - 15 A classical statement of the correspondence theory of truth is found in Aristotle's *Metaphysics*. For a study of the correspondence view and a defense of realism concerning truth in general, see William P. Alston, *A Realist Conception of Truth* (Ithaca, NY: Cornell University Press, 1996); and for a wide-ranging scholarly treatment of the theory of truth, see Richard Kirkham, *Theories of Truth* (Cambridge, MA: MIT Press, 1992).
  - 16 The schema in question is commonly called (Alfred) Tarski's *T schema* or the *disquotation principle*, as in the way illustrated it provides for eliminating quotation marks. For minimally technical discussions of its bearing on our understanding of truth see W.V. Quine, *Pursuit of Truth*, revised edn (Cambridge, MA: Harvard University Press, 1992), chapter 5; Mark Sainsbury, 'Philosophical Logic', in A.C. Grayling (ed.), *Philosophy: A Guide through the Subject* (Oxford: Oxford University Press, 1995); and Kirkham, *Theories of Truth*.
  - 17 The literature on truth, even on any one account of it, is immense. For a brief defense of the redundancy view see Quine, *Pursuit of Truth*, and for brief criticism of it see Alston, *A Realist Conception of Truth*. One might, like Quine in many places, speak of (declarative) sentences as true or false, and the minimalist and redundancy approaches have been taken to apply primarily to sentences. For reasons I cannot detail, so construing them may have some advantages over the standard terminology in the text, but it also raises serious problems.
  - 18 Although G.W.F. Hegel (1770–1831) is often regarded as the leading historical proponent of the coherence theory of truth, a clearer statement is probably to be found in Brand Blanshard, *The Nature of Thought*, vol. 2 (London: Allen & Unwin, 1940). Cf. Donald Davidson, 'A Coherence Theory of Truth and Knowledge', cited in Chapter 9.
  - 19 William James argued that truth is what works. See esp. his *Pragmatism: A New Name for Some Old Ways of Thinking* (New York: Longman, Green, 1907). Cf. John Dewey, 'Propositions, Warranted Assertibility, and Truth', *Journal of Philosophy* 38 (1941), 169–86.
  - 20 I have not presented these conceptions of knowledge and justification as analyses, in part because of how much they leave unspecified. But it may be argued that they do constitute analyses and indeed say as much

as an analysis, as opposed to a full-blown theory meant to foreclose options an analysis should leave open, ought to say. For a case to this effect regarding my proposed conception of knowledge, see James E. Taylor, 'Conceptual Analysis and the Essence of Knowledge', *American Philosophical Quarterly* 30 (1993), 15–26.



## 12 Scientific, moral, and religious knowledge

- **Scientific knowledge**

- The focus and grounding of scientific knowledge*
- Scientific imagination and inference to the best explanation*
- The role of deduction in scientific practice*
- Fallibilism and approximation in science*
- Scientific knowledge and social epistemology*
- Social knowledge and the idea of a scientific community*

- **Moral knowledge**

- Relativism and noncognitivism*
- Preliminary appraisal of relativist and noncognitivist views*
- Moral versus “factual” beliefs*
- Ethical intuitionism*
- Kantian rationalism in moral epistemology*
- Utilitarian empiricism in moral epistemology*
- Kantian and utilitarian moral epistemologies compared*

- **Religious knowledge**

- Evidentialism versus experientialism*
- The perceptual analogy and the possibility of direct theistic knowledge*
- Problems confronting the experientialist approach*
- Justification and rationality, faith and reason*
- Acceptance, presumption, and faith*



## 12 Scientific, moral, and religious knowledge

In perceiving the world around us we constantly acquire knowledge: of colors and shapes, objects and events, people and their doings. We also acquire knowledge as we look into our own consciousness. By thinking about things we already know and by drawing inferences from those propositions, we extend some of our knowledge. And through memory, we retain much of our knowledge. Justification is acquired, extended, and retained in much the same way.

But how far does our knowledge extend? We have explored how knowledge is transmitted once we have some, but not just how much we have in the first place or the range of subjects and questions to which it extends. I want to explore three important domains in which we are widely thought to have knowledge of truths that are central in guiding our lives: the scientific, the moral, and the religious. We should consider how knowledge and justification may occur in these domains. The task is immense; here I will simply show how the framework laid out so far can clarify knowledge and justification in relation to some important aspects of science, ethics, and religion. The focus will be more on knowledge than on justification. But much of what we learn about knowledge will apply to justification, and some of it may hold in other domains, such as that of art, history, or literature.

### Scientific knowledge

If we knew nothing through perception, we would have no scientific knowledge—roughly, the kind acquired by scientists in their professional investigations. And however much scientific knowledge one acquires by instruction and testimony from someone else, the discoveries that this knowledge represents must be made partly on the basis of perception. Perception is obviously crucial for laboratory work and observations of nature. Scientific work done without such first-hand experience, for instance by theorizing about nature or by mathematizing current information, still depends, directly or indirectly, on someone's perceptions.<sup>1</sup>

### *The focus and grounding of scientific knowledge*

How does scientific discovery fit into the framework I have developed? If we start with the idea that perception is basic for scientific knowledge, the picture which readily comes to mind is that one makes observations, inductively generalizes from them, and, through the inductive transmission of knowledge from one's premises to one's conclusion, comes to know the truth of a generalization. Imagine Galileo rolling balls down his famous inclined plane. He measures their acceleration, collects the individual items of knowledge he thereby acquires, arrays them as premises, and generalizes (in a special way) to his formula (Galileo's law), which gives the rate of acceleration for such balls in general. What does a case like this show?

First, the example rightly suggests that scientists tend to be interested in the nature and the behavior of kinds of things, such as accelerating objects, and that what is typically considered scientific knowledge is of generalizations: for instance, propositions about all freely falling bodies, not about any particular one. Knowledge of particulars is needed to obtain such general knowledge, but the former may be just ordinary perceptual knowledge. Granted, knowledge which is of a particular thing, but derived from a scientific generalization, say knowledge that a parachutist will land at 8 a.m., is scientific in the sense that it is scientifically based. Still, it is not the sort of knowledge regarded as paradigmatically scientific, or the kind scientists directly seek in trying to understand nature.

The second point suggested by the example is that scientific knowledge is inductively, not deductively, grounded. For instance, the generalization Galileo discovered concerning acceleration does not follow from the premises he formulated in expressing his data, say that ball 1 accelerated at a certain rate, that ball 2 accelerated at that rate, and so on. The generalization is strongly confirmed by such premises, but not entailed by them. Because it is not entailed by them, regardless of how well they justify it, they do not prove it. *Proof* of a proposition requires either a premise that entails it, or—in one ordinary use of the term—at least an event whose occurrence establishes it, as when we prove someone is in town by introducing him. In the former case, a proof is a conclusive deductive argument; in the latter it is a kind of active showing. Call the first proof in the argumental sense, the second proof in the behavioral sense.

Given the non-deductive (non-entailment) relation that holds between premises of scientific reasoning that yields, from knowledge of data, knowledge of a generalization, it is best to avoid calling the reasoning that supports a scientific generalization "scientific proof," as some do. The reasoning is not even deductively valid, much less the kind exemplified by a geometrical proof of a theorem from axioms.<sup>2</sup>

A third aspect of this case, however, may mislead. The example portrays Galileo simply observing and then generalizing, yet says nothing about why

he is observing. But he made his observations for a reason. This is to be expected; scientific knowledge normally does not arise simply from haphazard observations. Normally there is a question, such as whether falling objects speed up, that leads to observing a particular kind of thing.

Moreover, there is normally a tentative answer to such a question—a hypothesis—which both guides observation and sets the epistemic goal of the observations or the experiments that lead to scientific discovery. For instance, one might hypothesize that balls rolling down a plane speed up 100 percent in a given time interval and then observe their speed at each such interval to see if the initial speed doubles, quadruples, and so on. The goal is to show that the hypothesis is true and thereby come to know it, or, if it is not true, to find a hypothesis that does account for the behavior of the balls.<sup>3</sup>

A different example will bring out more clearly how both formation and testing of scientific hypotheses depends on imagination, observation, and antecedent knowledge. It has been observed that ants can find their nests after foraging, even in the desert where winds blow away any olfactory or visual traces their paths may leave.<sup>4</sup> Antecedent scientific knowledge indicated that ants have a neural capacity to determine direction by the position of the sun, but this would explain only how they return in the right *direction*. How can they “measure” distance? It was hypothesized that they might have a built-in pedometer by which they “counted” steps. Thinking of this required imagination, but even more imagination is needed to test the hypothesis. Suppose we could shorten the legs of some and lengthen the legs of others. Then we could reasonably predict that, in trying to find their nest after foraging, a significant number of the former would fall short and a significant number of the latter would overshoot. An ant population was partitioned into three groups and sent to forage. The first was normal, the second had legs cut off at the knee, and the third had them lengthened by attaching boar’s hair. The results were as predicted. This confirms (without *proving*) the pedometer hypothesis.

The central point here is that scientific knowledge does not automatically arise as we observe our surroundings. Normally, we must first raise questions about the world; in the context of background knowledge—which may come from previous science or everyday observation or both—they direct our inquiry. Only in the light of such questions are we in a good position to formulate hypotheses. These, in turn, are the raw material of scientific knowledge. Some are rejected, some are confirmed, and some that are confirmed may become known.

### *Scientific imagination and inference to the best explanation*

Scientific knowledge does not develop, then, simply by inferentially extending what we already know. Normally, it emerges only through using imagination, both in formulating questions and in framing hypotheses to answer them. This is one place where scientific invention occurs. It is not only

machines and devices that are invented but also hypotheses and theories. Invention and discovery may, however, coincide: if you invent a theory that is true, you may be said to discover the truth it states.<sup>5</sup>

The essential place of imagination in developing scientific knowledge is also illustrated by discoveries that result not from coming to know a generalization, but from apparent refutations of a proposition thought to be already known. The planet Neptune was discovered because the observed orbit of Uranus (the planet nearest to it on the Earth's side) was not as expected according to the laws of planetary motion, the principles astronomers used in describing the motions and paths of the planets. Partly to avoid having to revise well-confirmed laws, it was hypothesized that the deviation of Uranus from its expected orbit was caused by the gravitational effect of a more distant planet. The observations made to test this hypothesis revealed Neptune.

If the Neptune hypothesis was considered the best available explanation of the data, we could speak here of an *inference to the best explanation*, an abductive inference: roughly, an inference to a hypothesis on the ground that it best explains one or more other propositions taken to be known or justifiably believed—or at least taken to need explanation and to be candidates for justified belief or knowledge if the hypothesis turns out to be true. If two or more hypotheses are equally good explanations, we may justifiably choose between them as we see fit, though simplicity is generally viewed as a major consideration in making a rational choice in such cases—if indeed it is not taken to be an element in a good explanation in the first place.<sup>6</sup>

Once again, through the use of imagination, a hypothesis is formulated, and, through testing it, a discovery is made and new knowledge acquired. And again, the basis of the new knowledge is inductive, though unlike Galileo's knowledge about freely falling bodies, it is not a result of generalization. The pattern here is a successful case of inference to the best explanation. One imaginatively hypothesizes a gravitational influence by another planet as best explaining the deviation, tentatively infers that there is such a planet, tests the hypothesis, and, through positive results of the test, comes to know that the hypothesis is true.

We have already seen, in discussing the structure of knowledge and justification, that a proposition's explaining one or more others can count toward its justification, and that this point can be accounted for in either or both of two ways. First, one might note the role of explanation in increasing the coherence of the patterns to which the explaining and explained propositions belong. Second, we might take, as a principle of the transmission of justification from justified premises to a conclusion drawn from them, that *if we are justified in believing the premise that a proposition explains one or more others, then we tend to have some justification for believing the proposition itself*. The point here, however, is not mainly about justification but about discovery. We discover a great deal by seeking explanations and positing one or another hypothesis to explain the puzzling data. What we discover in this way commonly constitutes knowledge.

*The role of deduction in scientific practice*

These examples do not imply that deduction has no substantial role in the development of scientific knowledge. Far from it. Once we have a hypothesis, we typically need (valid) deduction to determine how to test it. For instance, one needs deductive mathematical reasoning to predict where to look for an as yet undiscovered planet, given a certain deviation in the orbit of Uranus.

Moreover, from very general laws, such as Newton's laws of motion, one may deduce less general laws, for instance the laws of planetary motion and Galileo's law of acceleration. (Actually, the best that one may be able to do is deduce generalizations which these laws only approximate; e.g., Galileo's law, which represents increase in acceleration as uniform, does not take account of slight changes in acceleration that are due to minute increases in gravitational attraction as the falling object nears the Earth. But this deduction still helps to explain why we should get approximately the results we do in testing or applying that law.) Deduction may, then, not only take us from a hypothesis and auxiliary assumptions, such as propositions about conditions of observation and the power of our telescope, to a proposition about a single event, say the sighting of a planet. Deduction may also take us from general laws or wide theoretical principles to less general laws or narrower principles. There is virtually no limit to the number of deductions we may make in search of new hypotheses.

Deductions of the second kind help to unify scientific knowledge. For example, they enable us to exhibit all the special laws of motion—for planets, for falling bodies, for projectiles, and so on—as instances of the general laws of motion. Even the behavior of gases, conceived as collections of molecular particles, can be explained by appeal to the general laws of motion. Their pressure in a container of air, for instance, is explainable in terms of how hard the particles hit its walls. This pressure, in turn, is connected with their temperature viewed as explainable by their average speed of movement. Thus, the laws of motion provide an understanding of what determines both pressure and temperature and indeed a model for visualizing their relationship. They also give us, by appropriate deductive inferences, a subsidiary law (Boyle's law) correlating the temperature of a gas with pressure at a constant volume.

Can we, then, have scientific proof after all, in which we validly deduce a special law of, say, motion, from more general ones? This is doubtful; for even if we might prove the special law relative to the more general ones, our knowledge of the more general laws is ultimately inductive and in a certain way inconclusive. That knowledge is based on inference to the best explanation or on generalization from observed data or, more likely, on a combination of these procedures. It might seem that we could at least speak of "relative proof" of scientific laws. But a proof in the unqualified sense must decisively establish what it proves; and as the generalizations that are our premises are apparently not decisively established, it is at best misleading to say that what we deduce from them is. Even a valid deduction from a well-confirmed true proposition is not necessarily a proof.

If our scientific premises in a deduction of one law from a more comprehensive one is not proved, and if indeed such premises stand to be revised as new discoveries are made—a common fate of generalizations in science—then what we know only through deduction from those premises is not proved either. These premises are indeed exposed to possible disconfirmation through the discovery of counter-evidence to what we deduce from them. If what is deducible from a set of propositions turns out to be false, then the premises that entail it are false as well: as presumptive guarantors of its truth, they must share in its falsification.

### *Fallibilism and approximation in science*

So far, I have sketched some of the ways in which what we call scientific knowledge develops, and I have criticized certain stereotypes. Science is not, for instance, a domain in which hypotheses are proved conclusively. Nor are they typically discovered by simply generalizing from observations we happen to collect. These points, however, imply nothing about whether scientific generalizations are true, or can be known.

If a common fate of generalizations in science is their eventual revision, should we reject the idea that there is scientific knowledge at all? Even the incomparable Sir Isaac Newton, as he was called, was shown to be mistaken on some important points. Even if discovering this took centuries, is there good reason to believe that any other scientific generalizations are, strictly speaking, true, in the sense that they describe the world both correctly and timelessly, and apply to past, present, and future? If some are true, that may not be typical. Commonly, what we call scientific knowledge is regarded by scientists as needing refinement and as possibly mistaken. Quite properly, their attitude is fallibilistic.

If scientists accept fallibilism regarding scientific beliefs—the view that these beliefs may be mistaken and the accompanying rejection of dogmatic attitudes—they nonetheless tend to hold a kind of objectivism: the view that there is an objective method for ascertaining whether beliefs about the world are true, that is (roughly speaking), a method which can be used by any competent investigator and tends to yield the same results when properly applied by different competent investigators to the same problem. Scientific method is widely taken by scientists and philosophers alike to be a paradigm of an objective method.

Since we cannot know propositions that are not true, one might conclude that we should really not speak of scientific knowledge at all, but only of relatively well-confirmed scientific hypotheses. This is a defensible position. We may prefer, however, to account for the apparent facts in a way that allows us to maintain that there is scientific knowledge. One possibility is that in speaking of scientific knowledge we are often speaking a bit loosely of what might be called approximate knowledge: well-grounded belief which holds true up to a certain level of precision in measurement—apart from “minor

inaccuracies,” we might say. Newton’s laws have not, after all, been found completely inaccurate. In building bridges, as opposed to dealing with astronomical distances or elementary particles, they seem an adequate guide, and their being only approximately true need cause no trouble in such practical applications.

One can insist that what is not precisely true is not known. But we could also say that what is approximately true in the scientific domain may be an object of approximate knowledge, and that beliefs of such propositions both are fallible and should be held with an openness to their revision in the light of new discoveries. I prefer the latter way of speaking. Why must we say it is false, rather than approximately true, that the circumference of a circle is its diameter times 3.1416, simply because pi can be worked out so much further? Indeed, given that pi can be carried out infinitely, how could we ever truly state the circumference of a circle if only absolutely precise propositions can be true?

There is, however, a second way to account for the apparent falsity of certain scientific generalizations. It seems that often their formulations are not properly taken to be absolutely precise, and that, rightly interpreted, they are true within the appropriate limits. Consider the general law that metals are conductors of electricity. Perhaps this should be interpreted with the understanding that certain abnormal (or for practical purposes impossible) conditions do not obtain. If metals should fail to conduct electricity at absolute zero, would this show the generalization false or simply that its appropriate scope of application is limited? The latter view seems more plausible.

These points in defense of scientific generalizations against the charge of wholesale falsity do not imply that none of those generalizations can be shown to be simply false. The point is that in some cases, instead of saying that scientific generalizations are not really true and hence do not represent genuine knowledge, it is preferable to speak either of approximate knowledge of a precisely formulated, but only approximately true, generalization or, as in this case, of unqualified knowledge of an imprecisely formulated truth. The difference is roughly that between approximate knowledge and knowledge of an approximation. In practice, however, there may be no easy way to decide which, if either, of these cases one is confronted with, or which indicates the better way to represent the state of one’s knowledge in a given scientific area.

### *Scientific knowledge and social epistemology*

I have so far spoken mainly about knowledge as individual belief. There is, however, scientific knowledge in journals no one entirely remembers. This can be called virtual knowledge as, although no one has it, many of us can easily acquire it. It is as accessible as our connections to our libraries; and a day could come when much information of that kind is more readily accessible by computer than are items of information we must carefully draw from long-term memory.



A natural assumption, however, even for such knowledge in the public domain, is that individuals have generated it and that only one or more individuals can acquire it. This is not an uncontroversial assumption. Science is often said to be a social enterprise, and some thinkers, including some proponents of feminist epistemology, hold that science is often unwarrantedly represented as “individualistic” and even atomistic.<sup>7</sup> One writer in feminist epistemology maintains that scientific knowledge is “produced by cognitive processes that are fundamentally social” and even goes so far as to say that scientific observation is “dialogic in nature.”<sup>8</sup>

Some feminist epistemologists, moreover, have held that “the gender of the inquirer influences the character of knowledge itself,” and “some people claim that women have gender-typical ‘ways of knowing’.”<sup>9</sup> But the last two claims are not widely defended by feminist epistemologists. They are independent of the view that scientific knowledge is social. They are not implied, moreover, by the thesis that the gender of a researcher may affect what is actually known, for instance because in survey research “subjects give different answers to questions depending on the perceived gender of the interviewer.”<sup>10</sup> This point is important for gathering knowledge, but it does not contribute to understanding the nature of knowledge or what constitutes evidence or justification.

Issues of these sorts are concerns of a subfield of epistemology that overlaps feminist epistemology at many points: social epistemology. If individual epistemology is roughly the theory of knowledge and justification as applied to individual persons—the enterprise we have been mainly engaged in—social epistemology is roughly the theory of knowledge and justification as applied to groups of (at least two) people.<sup>11</sup> The field of epistemology has traditionally focused mainly on individuals taken one at a time, with the notable exception of the status of testimony as a source of grounds for belief. Should there be this individual emphasis?

One answer is this. It seems quite possible that a single individual could have both knowledge and justification even if no group does, whereas it is not possible that there be a group that has actual knowledge or justification when no individual member of that group has knowledge of or justification for the proposition in question. *We* cannot know anything unless it is known by you or me or some other individual—though there are things we cannot *learn* by ourselves. Thus, in the case of a map of the world, knowledge by many people is needed to build the resulting social representation of geographical knowledge. Nonetheless, one of us could survive the rest and retain knowledge, or an omnipotent God could have created just one finite person with the capacity for knowledge and an environment in which the capacity is successfully exercised. In this respect, individual knowledge and justification are apparently logically prior to their social counterparts: the former is possible without the latter, but not conversely.

Virtual knowledge is an exception: “we” can have it in our libraries, though no one in particular has it. There is still another, quite different case in which



we can have virtual knowledge of something that no individual literally knows. You and I might each know one of two things that obviously imply a third, and, if we work together, we might thus be said to “know” the third but to have not quite articulated it. Perhaps I know that the assailant wore a full-length cape and you know the identity of the only person near the crime scene who wore one; between us we have a solution to the crime. This might be considered a kind of unrealized social knowledge.

### *Social knowledge and the idea of a scientific community*

Both examples of virtual knowledge are instances of what might be called social knowledge. But the first kind of virtual knowledge (illustrated by the library case), though only socially accessible, is individually realizable, whereas the second kind of virtual knowledge (illustrated by the detective case) is in a sense socially constituted. Any competent individual can get the former; only cooperating individuals can get the latter. One concern of feminist epistemology has been to emphasize the role of cooperation, not only in generating knowledge but also in characterizing human knowledge conceived in the kind of overall way that encompasses both individual and social instances.

The priority of individual over social knowledge by no means implies that given an adequate understanding of individual epistemology one can easily construct a social epistemology. That is not so. It is an interesting and difficult question what relation must hold between a group and its members in order for its knowledge or justification regarding a proposition to reside in one or more of those members. This brings us to a related kind of priority of the individual over the social.

As the metaphor of residing suggests, what a group knows or justifiedly believes is constituted by what one or more members knows or justifiedly believes. If we (human beings) know that wet grass is slippery, for instance, then some of us must have appropriately grounded true beliefs of that proposition (I leave aside unrealized social knowledge here). Not all of us have to; but if only a few of us do, then it would be wrong to say that we know, as opposed to, for instance, “It is known,” referring to the group as the context in which the knowledge occurs.

When a subgroup is intended by ‘we’, the situation differs. If the ‘we’ designates the scientific community, then it is permissible that only a few members know in order for the group to know. It can be true that “we” now know the mass of a proton even if only a very few have the appropriate information. This may be precisely because scientific knowledge is social in the sense of socially sharable (in a way introspective knowledge of one’s own sensations is not sharable). It is also both publicly accessible, at least in the sense that it is normally open to public testing and proper use by any competent investigator, and, typically, cooperatively generated, in that most of it arises from team efforts.

These three points about the social character of scientific knowledge in part explain what it is to say that there is a scientific community. It is arguable, moreover, that some kinds of knowledge are (empirically) impossible without cooperation, as in the case of knowledge of a theory whose development and confirmation require a team effort.<sup>12</sup> Indeed, a measure of scientific cooperation is commonly required even to maintain scientific knowledge once it is acquired. For given the problems and challenges facing scientific hypotheses and theories as new information is acquired, the grounds on which they are maintained will often be eroded unless new evidences or arguments are found to explain away new data found by opposing theorists or new investigations.

What the relevant examples of social knowledge seem to suggest is, on the one hand, the genetic priority of the social over the individual in the development of our scientific knowledge—with testimony as well as scientific cooperation playing a crucial role in producing that knowledge—and the constitutive priority of the individual over the social in epistemological matters: social knowledge, justification, and indeed belief are constituted by individual knowledge, justification, and belief respectively. The constitutive priority in question applies both to actual knowledge and justification, as I have illustrated, and (with such qualifications as are indicated by the detective case that illustrated unrealized knowledge) to virtual knowledge (and virtual justification).

The notion of virtual knowledge is implicit in the idea of “scientific knowledge” as the scientifically grounded “knowledge” accessible to us within a certain degree of readiness—for instance in widely available journals—but not residing in any person’s belief. It is because there is no belief of the relevant propositions that there are scare quotes around ‘knowledge’. What is not believed (or in any way in someone’s mind) is not literally known. Nonetheless, the constitutive priority of the individual clearly applies to virtual knowledge: plainly, we would not have access to the relevant knowledge unless some individual(s) among us did, at least in the sense of having access to crucial premises, as in the detective case.

Even if we do not consider individual knowledge logically prior to social knowledge, we still have the question of the status of apparent scientific knowledge (and certain other social knowledge). If we view science historically and do not idealize, it turns out that there is no unqualified answer to the question of whether what is called scientific knowledge is knowledge as conceived in this book. If we assume that there are some scientific propositions which are strictly true—and I see no cogent reason to doubt that there are some—then we apparently have no good ground for thinking that they cannot be known (or at least justifiedly believed). But the history of science indicates much change and extensive, apparently ceaseless correction of previously accepted hypotheses. For all its progress, it also gives us cause to wonder whether even at this advanced stage in scientific development we

grasp many scientific truths that future investigation will never show to be inaccurate.

I am inclined to say that in spite of both scientific error and the fallibility of scientific attitudes, we do have much scientific knowledge, even if it is all only approximate knowledge, or knowledge of approximations. But even if we have a great deal of scientific knowledge, if much of it is approximate or is knowledge of approximations, we are quite some distance from the artificial picture one might have of scientific knowledge as a set of beliefs of precisely formulated and strictly true generalizations, arrived at by inductive transmission of knowledge from its basic sources in experience and reason. Those sources remain basic, and scientific method provides an objective way of building on them. But there is no straightforward transmission, or, when transmission occurs, any clearly final destination toward which it must proceed.

### **Moral knowledge**

The possibility of moral knowledge raises rather different sorts of questions from those just explored. Moreover, whereas there is a widespread tendency to take for granted that there is much scientific knowledge, there is a widespread inclination to take moral judgments to be at best culturally conditioned assertions with no claim to genuine truth.

Consider the judgment that cruelty to children is wrong. A clear application of this principle might be the more specific judgment that it is wrong to thrash a two-year-old for accidentally spilling milk.<sup>13</sup> We accept this, but do we know it? Suppose someone denies it or simply asks us to justify it. It does not appear that we can establish it scientifically. It is apparently not a scientific judgment in the first place. Furthermore, it is not in any obvious way grounded in perception; and if it is grounded in reason, it does not seem to be so in the straightforward way the representative self-evident truths discussed in Chapters 5 and 6 apparently are. Many find it natural to consider this judgment to be grounded in our culture and to be accepted simply as part of the social fabric that holds our lives together. It would then be a socially accepted judgment but would not express social knowledge.

### ***Relativism and noncognitivism***

There are at least two major variants of the view that moral judgments are somehow grounded in our culture. One of them allows that they are true, but only in a qualified sense that reflects their being tied to the culture in which they occur. The other takes them not to belong to truth-stating discourse at all. Let us consider these views in turn.

The first view—a kind of *relativism*—says roughly that moral judgments are true relative to our culture (or even some subculture), but not unqualifiedly true, as judgments of fact, such as that a green field lies before me, may

be. This is not to say merely that moral judgments are relativized—true in relation to certain circumstances, as when ‘We should not kill’ is said to be true relative to non-self-defensive conditions. This relativization view does little more than restrict moral judgments, often uncontroversially, to the circumstances in which they properly apply. Virtually all writers in ethics are “relativizationists.” Relativism, by contrast, is the thesis that there are, say, American moral truths, British moral truths, Chinese moral truths, Danish moral truths, and so on, but no universally valid moral truths (or standards)—even if some moral principles are universally held.

The second view that (in a way) takes moral truths to be culturally grounded—the attitudinal view, also called *expressivist*—says roughly that such judgments are not literally true at all; rather, they are expressions of moral attitudes, not assertions of a proposition, hence not “cognitive.” Normally, these are attitudes rooted in the culture of the person judging, but a noncognitivist might allow a subgroup or even an individual to qualify as having a morality in the sense relevant to providing the basis of moral judgment.

On one version of the attitudinal view, to say that cruelty to children is wrong is like uttering ‘Cruelty to children!’ in a tone expressing revulsion and adding, ‘I (hereby) condemn it!’ The attitude is moral both because of what it concerns and because it represents a certain kind of cultural stance.

Such attitudes may be reasonable or unreasonable and may be defended with reference to what is true or false; but attitudes are not themselves true or false.

On both views there is no moral knowledge, since either there are no moral propositions at all (the negative claim of the attitudinal view), or there are at least none that are true or false unqualifiedly (the negative claim of the relativist view in question), whereas propositions expressing empirical or a priori truths may be unqualifiedly true or false. The attitudinal view is thus committed to (ethical) noncognitivism: roughly the claim that there are no moral propositions to be known, or otherwise “cognized.” The (ethical) relativist position in question need not endorse noncognitivism. But doing so can clarify a relativist view, because, as our examples will show, it is not clear precisely how propositions can be true in a way other than the way in which empirical and a priori truths are, and it is clear (at least in outline) how moral attitudes can be rooted in a culture.

The attitudinal interpretation of moral judgments is on the surface the more radical view. It implies that there simply are no moral truths—or moral falsehoods either. There are no moral propositions to be known, or to be justifiedly or even mistakenly believed in the first place. What makes this apparently radical view plausible?

Suppose one is very impressed with the basic sources of knowledge as our only routes to knowledge, and one notes that apparently no propositions known on the basis of sense experience seem to entail the truth of any moral judgment. For instance, even the proposition that cruelty to children causes

them pain does not entail that cruelty is wrong. Needed surgery, after all, may cause them pain yet not be wrong.

When we judge something to be wrong we apparently go beyond the evidence of the senses, and indeed beyond scientific evidence. For example, suppose that (as is surely possible) we know scientifically that in fact cruelty to children commonly breeds brutality in its victims. Unless we already know or are justified in believing that breeding brutality is wrong, the fact that cruelty to children breeds brutality does not justify us in believing that cruelty to children is wrong.

It seems to noncognitivists, then, that we cannot know that cruelty to children is wrong just on the basis of its causing brutality; this fact would (deductively) ground that knowledge for us only if we already knew that brutality is wrong. Now suppose we also assume that nothing known *a priori* entails that cruelty to children is wrong: no logical truth, and not even a synthetic proposition like the truth that nothing is red and green all over at once. These points serve as premises for the negative conclusion that there is no moral knowledge. For if knowledge is grounded in the basic sources and moral judgments are not grounded in them, then moral judgments do not constitute knowledge.

There is also a positive thesis held by the attitudinal theory: that (even though moral judgments do not express propositions) moral judgments do express significant attitudes. A main reason for saying they express attitudes is that we are not neutral in making moral judgments; we are (normally) pro or con regarding what we morally judge. Normally, we at least indirectly commend or condemn when we make a positive or negative moral judgment.

Now suppose we combine the positive view that moral judgments express attitudes with the negative view that there are no moral propositions, which itself implies that there is no moral knowledge. It is now plausible to conclude that the point of making a moral judgment is not to assert something but to express a positive or negative attitude and thereby to influence human conduct, if only by endorsing or condemning one or another kind of behavior. Many noncognitivists hold this third position.

Noncognitivism must not be taken to imply that in moral matters "anything goes." Noncognitivists may grant that moral attitudes may be unreasonable, say based on misinformation or prejudice. The view can thus allow that there are moral mistakes. But mistakes that are specifically moral are mistakes in attitude, not about what is true or false.

The relativistic view that moral judgments are culturally grounded endorses the first argument just set out, based on the premise that those judgments are not anchored in the basic sources of knowledge, but not necessarily the second, attitudinal argument. On this relativist view, although moral judgments are not rooted in those basic sources, they are learned as we absorb our culture, and they may thus share with judgments that do represent knowledge a wide social acceptability. Still, we are at best entitled to expect them to be accepted within our society (or one that morally agrees with it), and they

are at best “true for” one or another society. They are not unqualifiedly true, hence not genuinely known in the sense that implies cross-culturally valid standards of evidence. If they express propositions, those propositions are assertable in our culture, but not unqualifiedly true.

### *Preliminary appraisal of relativist and noncognitivist views*

To assess the relativist and noncognitivist views, consider first the part of each position not shared with the other one. Let us start with the attitudinal aspect of noncognitivism and proceed to the cultural groundedness thesis characteristic of relativism.

Is there an alternative explanation of the attitudinal aspect of moral judgments? Might they be true or false and still have, for instance, the commendatory or condemnatory force they do? Take the utterance, ‘The curtains are on fire!’ If sincerely uttered by any normal person, this would commonly express alarm and be meant to evoke action. But it is clearly factual, and it can be unqualifiedly true.

Moreover, it seems to be because of its factual content that the judgment that the curtains are on fire expresses the kind of alarm it does. Perhaps certain statements of those facts that are significantly and obviously linked to human concerns are no more attitudinally neutral than are typical moral judgments. For this reason (among others), one might resist the idea that the only distinctive function or even the major function of moral judgments is to express attitudes, as opposed to asserting propositions. It might just be that the facts supporting moral judgments are socially important in a certain way, as is the curtains’ being on fire.

As to the point that moral judgments are, in some cases, culturally tinged and differ from one culture to another, there are at least two related ways to explain this. One concerns the social origins of their acceptance, the other the content of such judgments.

First, cultures differ in what they accept as a moral judgment. Whereas contemporary Western societies might construe binding the feet of young girls to keep them from becoming large and “indelicate” as cruelty, the practice was not considered wrong by at least certain segments of Chinese society in the past. It might be replied that ‘cruel’ is a moral term, and hence there might be a moral disagreement after all—not about whether cruelty to children is wrong, but about whether this practice is cruel. Suppose there is disagreement about that. Why might it not be resolvable by appeal to non-moral facts? Disagreements about whether certain punishments are morally permissible may turn on whether they tend to deter crime. Here a moral disagreement about what counts as morally permissible punishment does turn on the parties’ view of relevant non-moral facts.

Second, we must distinguish—even if we must also interconnect—genetic and justificatory considerations. The beliefs we express in making moral judgments may be learned through absorbing a culture, even if what justifies

those beliefs or renders them knowledge does not depend on absorbing a culture. Clearly, the origin of a belief need not be what justifies it, nor need it reflect the truth of the proposition believed, if it is true. Thus, we might learn a moral principle through something characteristic of our culture (such as moral education), even though what justifies it is not grounded in our culture but, perhaps, something pertaining to human life as such. Our moral education might reflect this universality; but one could first learn a moral truth from an unreliable source—say, someone who deceitfully calls another person unfair—and later get a good justification for it.

Suppose that our moral beliefs do arise from our education and culture, and are in this sense culturally and historically “conditioned,” as it is sometimes put by theorists working in the sociology of knowledge, a field that overlaps both social epistemology and feminist epistemology. There is good reason to say that at least many of our scientific beliefs are also culturally and historically conditioned. If we need not thereby regard the relevant scientific beliefs as culturally relative, why should we so regard moral beliefs?

One might think that unless they are scientifically justified, moral judgments are merely true for those who hold them. But both moral and scientific judgments, moreover, are “true for” the social groups that hold them, at least in the sense that the people in question believe them. Does that not indicate a kind of relativity in both cases? That is doubtful: anything we believe is in that sense true for us—it is believed by us to be true. If this is how moral judgments are relative to those who make them, their “relativity” characterizes even simple self-evident truths. Isn’t it true for all of us that if the spruce is taller than the maple then the maple is shorter than the spruce?

Is there some other sense of ‘true for’ that discriminates between the sense in which anything we believe is true for us and the elusive sense in which moral judgments are, according to some relativists, true for some people and not others? One might try taking ‘true for’ as equivalent to ‘true from the (cultural) point of view’; but what does this come to beyond saying that ‘It is true for my culture’ means ‘My culture believes it’?

Another possibility is that ‘true for’ means something like ‘successfully works for’ or alternatively ‘conforms to the practices of’. This would be illustrated by ‘One ought to drive on the left side of the road’ is true for the British but not (many) others. But this does not yield a general relativism or at least not one that sets ethics apart from science. It is equivalent to something like ‘In Britain one ought to drive on the left’, and that simply specifies circumstances in which the judgment applies, just as ‘at sea level’ specifies when the air pressure on Earth is 14.7 pounds per square inch.

This circumstantial relativism—as we might call it—is simply the plausible view that what we ought to do depends on the circumstances we are in. It says nothing about the status or the nature of the truth of moral principles once they are stated in relation to—relativized to—the circumstances they apply to. It leaves open that they might then be seen to be true or false in the usual sense appropriate to propositions about the empirical world.



There are certainly different kinds of circumstances in different cultures, and there may be important moral principles true for one society and not another, in the sense that in one of the societies, but not the other, people generally believe them. But, as the analogy to scientific disagreement indicates, that would show nothing about whether moral principles or judgments are relative in any sense implying that they cannot be known or justifiedly believed.

### *Moral versus “factual” beliefs*

We are now approaching the heart of the issue concerning the possibility of moral knowledge. Recall the objection that scientific beliefs, but not moral judgments, are testable by experience and reason, say by appeal to perceptual facts and deductive logic, and hence moral judgments are relative in a way scientific judgments are not. We are back to the argument which the relativist and attitudinal views commonly share: that since experience and reason do not ground moral judgments, those judgments cannot express knowledge. This argument must be squarely met.

The first thing to be stressed is that from the premise that moral judgments are not formally deducible—derivable by the standards of deductive logic—from facts, it simply does not follow that they are not justifiable by appeal to facts. That this conclusion does not follow is evident from our discussion of scientific knowledge, which illustrates that knowledge can arise through inductive transmission from evidential premises. Scientific generalizations, for example, are inductively known on the basis of facts, such as observational data, which we use to confirm them. If there can be scientific knowledge on this basis, then there can be knowledge based on inductive grounds, grounds that do not entail any proposition known on the basis of them. Why, then, should there not be inductively grounded moral knowledge?<sup>14</sup>

One reply to this argument is that moral generalizations are not even inductively supported by the facts. But is that true? We certainly cite facts to justify moral judgments. I might justify my judgment that I ought to meet with Jane by citing the simple fact that I promised to. This does not prove that I ought to meet with her, but it surely provides a good reason for the judgment that I ought to. There is, moreover, a third possibility we should examine: that even if such a fact does not imply a moral judgment by the rules of logic, it implies it in a different kind of a priori way.

### *Ethical intuitionism*

This brings us to a major account of moral knowledge, one quite different from the deductive and inductive ones so far specified. Suppose someone asks why I should keep my promises in the first place. I could perhaps explain why I believe this. But suppose that I cannot justify it by appeal to anything more basic. This would not show that I do not know or justifiably believe it.



At some point or other in defending a perceptual judgment I may be equally unable to give a further justification. It would not follow that the judgment does not express knowledge or justified belief.

The issue before us should be explicitly considered in the light of what we saw concerning the structure of knowledge. A foundationalist may say that (with some special exceptions) the principle that one should keep one's promises, or at least some more general principle, such as that people should be treated with respect, is self-evident, hence knowable a priori, and needs no defense by derivation from prior principles. This intuitionism does not claim that everyone who considers the relevant principle finds it obvious; but that same point will hold for certain truths of logic which, when finally understood, are comfortably accepted as self-evident. The crucial thing is that, through intuitive reflection, the principle can be seen to be true.<sup>15</sup>

Chapter 5 treated intuition as a rational capacity: roughly, a non-inferential, apprehensional capacity by whose exercise what is intuitively believed or known is believed or known. But I have said little about *intuitions*, understood as a kind of exercise of that capacity. To understand intuitions we must consider at least three notions: (1) intuitions *that p*—*cognitive intuitions*; (2) a proposition's being *intuitive*—evoking (under certain conditions) what might be called *the sense of non-inferential credibility*; and (3) *objectual intuitions*, roughly direct apprehensions of either (a) a concept or (b) a property or relation, such as the property of being a promise, the property of being unjust, or the relation of entailment. (1) and (2) will be my main concern.

On one view, cognitive intuitions are a kind of belief. This would explain why we rarely speak of an intuition that *p* without presupposing that the cognition in question entails belief that *p*. Moreover, if we are asked whether *p* is true, then, if we are aware that we do not believe *p*, but it (non-inferentially) seems true, we are likely to say, in careful usage, just that—that it seems true or, sometimes, that it seems intuitive. If, however, cognitive attitudes are conceived broadly, as those with truth-valued objects, then what is not believed but, in the relevant phenomenal way, seems true to us may be considered an intuition of ours. For some philosophers this is the primary concept of an intuition.<sup>16</sup> We might thus divide cognitive intuitions into those that are *doxastic*, that is, a kind of belief, and those that are not but embody a disposition to believe. I prefer, however, to call the latter *intuitive seemings* rather than intuitions. This terminology is harmonious with the view of intuitions held by most who take seemings to be the primary (or only) cases of cognitive intuitions.

More important than which terminology we use is the relation between doxastic intuitions and intuitive seemings. An intuitive seeming that *p* can be an evidential ground for believing *p*, somewhat as a sensory seeming can be a ground for believing that the paper before one is white. Intuitionists have typically presupposed this, sometimes with the idea that just as a perceptual seeming—say, its visually seeming that there is paper here—is evidence for that proposition (indeed for a range of “observational” propositions), an intuitive seeming that *p* is evidence for *p*.

Foundationalists will tend to argue that such a response—considering *p* true on the basis of its being intuitive—is legitimate when we get to certain stages in a process of justification. For them, some beliefs (including many non-self-evident kinds) are foundational in a way that warrants holding them without doing so on the basis of prior premises. If that were not true, then we could not be justified in holding anything. This is not in the least to suggest that the foundationalist must cut off discussion or simply reassert the proposition in question. Clarification by examples may be adduced. Misguided objections may be refuted. An unexpected premise might be found after all. That is possible even for a self-evident proposition. Being evident in itself, such a proposition does not stand in need of a premise, but it is not necessarily incapable of being evidenced by one.

Coherentists will certainly be willing to go on arguing by appeal to propositions that provide support by appropriately cohering with the one in question. They might point out that if we do not keep our promises, life will be unbearable, and then, for each thesis attacked, defend it with respect to one or more others. The objector may not be pacified by this approach either. But neither approach can simply be rejected out of hand. To be warranted in rejecting either approach, one must have a plausible alternative conception of knowledge and justification. What would it be? That is far from evident, as we shall soon see in exploring skepticism.

These responses in support of the possibility of moral knowledge do not go as far as one might like. They rest on limited analogies and on simply showing that the case against moral knowledge is inconclusive. There are two other important responses we should consider. One, defended perhaps most powerfully by Kant and later Kantians, as well as by intuitionists, construes knowledge of moral principles as *a priori*. The other, defended perhaps most powerfully by Mill and later utilitarians, represents moral principles as empirical. In either case, moral knowledge and moral justification are grounded in the basic experiential and rational sources I have been discussing.

### *Kantian rationalism in moral epistemology*

To understand the first, broadly Kantian, response, consider another application of the principle that cruelty to children is wrong: the proposition that flogging infants for pleasure is wrong. There is some plausibility in saying that we know this. Intuitionists would tend to say we know it (or can know it) non-inferentially; Kantians would likely hold that we know it as an obvious application of Kant's famous categorical imperative, which, in one form, says that we are to act only on principles that we can (rationally) will to be universal laws of nature obeyed by us all.<sup>17</sup>

The proposition that flogging infants for pleasure is wrong seems plausible on even brief reflection about what it is to flog infants. It is difficult to conceive exceptions, and it is certainly difficult to conceive circumstances that would lead rational persons not to endorse it provided they are taking either the point of view of universalizability or the equally Kantian point

of view of commitment to treating people as ends in themselves and never merely as means. To be sure, perhaps such cruelty as flogging an infant could in special cases be excusable, even if not morally right. Terrorists might electrically manipulate my brain so that I change in personality and am somehow brought to flog an infant (perhaps even to do so for pleasure); but even then I would be doing something wrong, though in an excusable way.<sup>18</sup>

Consider another example, a modest version of something more powerful: we ought to treat people equally in matters of life and death, say in regard to wartime military service, unless they differ in some relevant way (and not merely in being different people). This is a kind of principle of consistency, not logical consistency, but something like consistency in roughly the sense of using a principled policy or procedure in making decisions. It says that such preferentially inconsistent treatment is *prima facie* wrong and that differential treatment in these mortal matters must be justifiable by a difference.

Granted, the principle does not specify what kind of difference is relevant, for example that the health of candidates for organ transplant is relevant and their skin color is not. Specifying relevant differences is a further step. But the principle is still a moral one, and it commits us to the important requirement that there be a reason to justify the indicated differences in treatment. Particularly since it is a kind of consistency principle, there is some reason to believe that if it is true, it is knowable *a priori*, though defending this idea would be a major task.

As this perspective on equal treatment suggests, it is natural to take the Kantian view to be internalist in its moral epistemology: it is by the use of reason, and hence through grounds accessible to reflection, that we know and can justifiably believe sound moral principles. We may need much experience to understand moral concepts; but once we understand them, sufficient reflection on them provides justification for basic principles of action and thereby for moral principles.

### *Utilitarian empiricism in moral epistemology*

The second response to relativist and noncognitivist views in moral epistemology, the response of Mill's utilitarianism, is very different. Its first point is that moral judgments are knowable on the basis of factual knowledge of how acting in accordance with them would contribute to producing something intrinsically good, hence good in itself, independently of what it leads to. Second, Mill maintained that only pleasure and freedom from pain are good in themselves. He apparently believed that if these two premises can themselves be known (as he thought they could be), they justify holding, as one's fundamental moral principle, something like this: that precisely those acts are right which contribute at least as favorably to pleasure (and freedom from pain) in the relevant population as any alternative available to the agent. (I leave aside the points Mill raises later about qualities of pleasure.)<sup>19</sup>

Because, on Mill's view, we can determine what these optimal acts are,

that is, what acts have the most *utility*, by a combination of common sense and scientific procedures, moral judgments are knowable in the same way as common sense and scientific statements. By contrast with Kant's moral epistemology, on which we have internal (a priori) access to the grounds of basic moral truths, Mill's moral epistemology (and that of at least the majority of utilitarians) is externalist. For them, we have access to grounds of moral truths only through considerations about the consequences of actions for pleasure and pain, and those considerations require observational or other kinds of inductive evidence. Given the central role of consequences in determining what is right, utilitarianism is a form of *consequentialism*. Consequentialism takes other forms, and most are epistemologically similar.

A question that now arises is how we know that pleasure or anything else is intrinsically good. Mill argued that (for one thing) we can know this by determining what people by nature actually desire for its own sake. But the utilitarian approach is by no means committed to that view (which many commentators on Mill find implausible). For instance, it might be argued instead that what is intrinsically good is what people want or would want for its own sake provided their wants are adequately rational, say held in the light of reflection that is logically and scientifically rational, vivid, and appropriately focused on the nature of the thing wanted.<sup>20</sup>

### *Kantian and utilitarian moral epistemologies compared*

The Kantian and the Millian utilitarian responses to challenges to moral knowledge are nicely parallel to Kant's and Mill's views of the truths of reason. On Kant's rationalistic view, moral principles are (synthetic) a priori. On Mill's empiricist view (only part of which I have stated), moral principles are empirical.

There is a further epistemologically interesting contrast here. On Kant's approach, or at least on some approaches of the same rationalistic kind, such as most versions of intuitionism, there can be direct (non-inferential) moral knowledge. For on these views some moral principles are basic in a sense implying that knowledge of them need not be inferentially grounded in knowledge of any other propositions. For utilitarianism, there cannot be direct moral knowledge except in special cases. The main and perhaps only cases seem to be these. First, some direct moral knowledge is only memorially direct, that is, direct as preserved in memory but originally indirect and now direct just by virtue of one's forgetting one's evidential grounds for it, as we forget the steps in proving a theorem and remember only the theorem. Second, some moral knowledge is testimonially direct, that is, non-inferentially grounded in testimony, where (for the utilitarian) this requires that at some time someone (say, the attester) knew the truth inferentially.

Both of these memorial and testimonial cases would be *secondary knowledge*, since the knowledge depends on other knowledge (of a non-testimonial, non-memorial kind) of the same proposition and is not primary in the way

that, say, perceptual knowledge is. Secondary knowledge need not, however, be inferential, since it need not be based on other knowledge. For Mill, knowledge that, say, keeping one's promises is obligatory would ultimately depend on someone's knowing a good deal about the effects of promise-keeping on happiness. We can know the principle through parental teaching in the course of our moral education. We could establish it for ourselves by studying human behavior and then retain our knowledge of it after forgetting our grounds. But no one could know it directly unless someone knew it inferentially, through evidence.

This difference between Kant and Mill—the former providing, as do intuitionists, for direct moral knowledge and the latter not—is no accident. Implicit in Mill's utilitarianism is the view that moral properties, such as being obligatory, are unlike sensory properties in not being directly experienced or otherwise directly apprehensible. As an empiricist, and thus experience-based, moral theory, it must treat knowledge of moral truths as ultimately indirect (unless, as has sometimes been done, it posits moral experience as a source of knowledge that grounds knowledge rather in the way perception does). Thus, even if, by memory, I have some direct moral knowledge, no moral knowledge is *independently basic*, in the sense that it need not at any time be grounded in another kind of knowledge. If I know that cruelty to children is wrong, it is by virtue of my (or someone's) knowing that it does not contribute optimally to happiness in the world.

For a broadly Kantian view, by contrast, we can rationally grasp this principle, at least as a consequence of a more general principle. For intuitionism, we can sometimes even directly grasp a moral principle, say that arbitrarily unequal treatment of persons in matters of life and death is wrong. On both these views, then, we can have moral knowledge which is direct and independently basic. But even if Kant is best interpreted as construing the most general moral knowledge as depending on non-moral premises, he took all fully general moral knowledge to be deductively derivable from (and only from) a priori premises and thus itself a priori (at least in the provability sense).

Should Kantian or utilitarian or intuitionist views convince us that there is moral knowledge? From what I have said about them here, it is not obvious that they should. But, when carefully developed, they are each plausible, and each may be held with the attitude of objectivistic fallibilism that is also appropriate to scientific views. Each view certainly seems to warrant the conclusion that there can be moral knowledge; and apparently there is some, despite the sorts of relativistic and attitudinal arguments I raised to indicate why some thinkers deny its possibility.

There are, of course, other issues that should be explored in deciding whether moral principles or judgments can be known, or even justifiably believed. There may, for instance, be sources of moral knowledge, such as a special moral faculty analogous to perception. But I see no compelling reason to believe there are. Suppose, however, that we do have a special moral faculty. Presumably, it is a kind of rational capacity whose insights are rational

ones capable of being known or at least justifiedly believed. Whatever the problems that remain, perhaps enough has been said to connect them with the epistemological framework developed in this book. Certainly, locating moral judgments in that framework is at least a good way to approach the question of whether they can constitute knowledge.

## Religious knowledge

The case of possible religious knowledge is different in many ways from that of possible moral knowledge, but it can also be clarified in the light of some of the concepts and principles introduced in this book. Again, I want to be brief and to start with the negative view, in this case that religious propositions are simply beyond the scope of human knowledge. I have in mind mainly propositions about God, such as that God exists, brought order out of chaos, created the universe, or loves us: propositions that are not merely religious in subject matter but also imply or presuppose that God (or some spiritual reality with a central place in a religion) exists.

Why would it be thought that no religious propositions are known? The most common basis for holding this view is probably much like the most common reason for holding that there is no moral knowledge, namely, that religious propositions cannot be known either *a priori* or on the basis of experience, say by inferring God's existence from the premise that God's designing the universe best explains the order we find in it.

Both aspects of this negative claim have been discussed by philosophers and theologians at great length, and there are well-known arguments for the existence of God meant to provide knowledge that God exists. Some of these use only *a priori* premises; others use only empirical premises. For instance, the ontological argument, in one form, proceeds from the *a priori* premises that God is supremely perfect (has all perfections in the highest degree), and that existence is a perfection, to the conclusion that God exists. By contrast, the argument from first cause (in one form) uses the empirical premise that there is motion, together with the general premise that there cannot be an infinite chain of causes of motion, and concludes that God, as an unmoved first mover, exists.

There is a vast literature about these and all the other historically important arguments for the existence of God.<sup>21</sup> I am not concerned here with arguments for God's existence. All I want to say about those arguments is that nothing in the framework I have developed implies either that there can or that there cannot be cogent arguments for God's existence. For instance, nothing said about the basic sources of knowledge or about its transmission implies that those sources could not in some way lead to arguments yielding knowledge of God or of some other spiritual reality. The same point applies to justification of beliefs about God or about some spiritual reality, and both points hold within either a foundationalist or a coherentist epistemology.

*Evidentialism versus experientialism*

But what about the possibility—less often discussed than arguments for God’s existence—of direct (non-inferential) knowledge of God? Does the framework of this book rule out that possibility? General epistemological considerations have sometimes been thought to do so, but they do not. Indeed, if there can be what I have called natural knowledge, as in the case of direct knowledge of arithmetical results ordinarily knowable only through lengthy calculation, then there is some reason to think that knowledge can be built into a person in such a way that the person could have direct knowledge of God. (The kind of knowledge in question was countenanced by John Calvin and its source dubbed the *sensus divinitatis*.)<sup>22</sup> To be sure, there may be less mystery about how a mere calculating mechanism could be built into the brain than about how knowledge of an external, spiritual reality could be. But a mystery is not an impossibility.

If, however, it is even possible that there is an all-powerful (omnipotent) God, then that God could create such direct theistic knowledge. If there can be such knowledge, then one form of what is called evidentialism is mistaken, namely, evidentialism about theistic knowledge, the view that knowledge of God is impossible except on the basis of adequate (propositional) evidence. On this view, religious experience, say as described by mystics, is not considered evidence; the kind intended is not the non-inferential “evidence of the senses,” such as we have for there being paper before us, but the sort ordinarily called evidence, which could be expressed in premises enlisted to support theistic conclusions.<sup>23</sup>

How might evidentialism apply to justification? Recall the *prima facie* cases of direct knowledge of something that is ordinarily knowable only through evidence or inference, such as the result of multiplying two three-digit numbers. If there is direct knowledge here, it need not be a case of justified belief. So we cannot use such examples to refute evidentialism about theistic justification: the view that justified beliefs about God are impossible except on the basis of propositional evidence.

Could one be directly justified in believing such religious propositions as that God exists? Would this require one’s having a sixth sense, or a mystical faculty? And even if there should be such a sense or faculty, would it generate justification directly, or only through one’s discovering adequately strong correlations between its deliverances and what is believed through reason or ordinary experience, for instance through one’s religious views enabling one to predict publicly observable events? In the latter case, the sense or faculty would not be a basic source of justification. Before it could justify the beliefs it produces, it would have to earn its justificational credentials through a sufficient proportion of those beliefs being confirmed through other sources, such as perception and introspection.

There is, however, a way to resist evidentialism and argue for the possibility of direct justification of certain religious beliefs without assuming that



there are any sources of justification beyond reason and normal experience (though normal experience need not be through the five senses). In particular, this approach need not posit either mystical apprehensions, such as overpowering, ineffable, otherworldly experiences, or special divine revelations, whether in those experiences or in the presence of apparently miraculous changes in the external world.

I call the position I have in mind *experientialism*, since it grounds the justification of some very important religious beliefs in experience rather than in evidential beliefs or direct rational apprehension. Religious people sometimes say that, in perfectly ordinary life, God speaks to them, they are aware of God in the beauty of nature, and they can feel God's presence. Descriptions of these sorts might be considered metaphorical. But if God is, as many think, properly conceived as a (divine) person, these avowals might have a literal meaning.

It is natural to object that all one directly hears in such experiences is a special kind of voice (presumably in one's mind's ear), that all one directly sees is the natural beauty which one takes to manifest God, and that one simply feels a spiritual tone in one's experience. From these moves it is easy to conclude that one is at best indirectly justified in believing one is experiencing God. After all, one believes it inferentially; for instance, on the basis of one's belief that the voice one hears is God's, one might believe that the beauty one sees is a manifestation of divine creation; and so forth.

### *The perceptual analogy and the possibility of direct theistic knowledge*

To assess the case just made to show that theistic beliefs are inferential and so not candidates to be directly justified (or direct knowledge), compare perception. Suppose it is argued that one is only indirectly justified in believing there is a green field before one, since one believes it on the basis of believing that there is grass, a green textured surface, and so on. Must we accept this? I think not. I do not normally even have these beliefs when I believe there is a green field before me, even if I do see it by seeing its grassy surface.

The matter is far more complicated than this, however. It may be argued that as God is both infinite and non-physical, one cannot be acquainted with God through experience. But this argument will not do. Even if a stream were infinitely long, I could still see it by seeing part of it. Seeing an infinite thing is not seeing its infinity.

But, if seeing an infinite stream is not seeing its infinity, then how can seeing it be a basis for knowing that the stream is infinite? Similarly, supposing God is experienced, how can the experience reveal that it is God who is experienced? The problem is not that God is non-physical. The non-physical can be quite readily experienced, and indeed in a direct way. Thus, my experience of my own thinking presumably need not be of something physical, even if in fact it is of something physical, say a brain process; and even if it



must be (because of some necessary connection that might hold between the mental and the physical), it is not experience of, say, my thoughts as physical.

The problem, then, is not that there cannot be experience, even non-mystical experience, of God. It is (in part) that if experiencing, say, God's speaking to one, is possible, it is not clear how one could know (or justifiedly believe) that it is God speaking. How would one know that one was not having a merely internal experience, such as talking to oneself in a voice one thinks is God's, or even hallucinating a divine voice? (Such skeptical questions will be discussed in the next chapter.)

In part, the question is how one might recognize God. Plainly, this requires having a concept of God. But that is acquirable without already having knowledge of God's existence. One also needs a concept of, for instance, a sonata to recognize one. These concepts are very different, but either one can be acquired without actually knowing of the existence of (or experiencing) what it represents.

Here it is important to recall the perceptual analogy. Why would it be less likely that my experience of looking toward the green field is hallucinatory? It is true that there is a difference: we can, with all the other senses, verify that we see a green field, whereas God may seem perceptually accessible at most to sight and hearing—presumably indirectly, as God is seen in appropriate things and heard through hearing voices, perhaps inner voices, that are not literally God's (at least if a being's voice must be physically grounded in a physical embodiment, though even in that case, some would argue that God's voice was physically embodied in Christ).

Even if God is accessible only to sight and hearing and, in any case, only indirectly, it does not follow that knowledge and belief about God are indirect. As we saw in exploring the sense-datum theory, we can know one thing through another without inferring facts about the first from facts about the second. Thus, the force of this difference between the possible perceptual accessibility of God and that of physical objects can be exaggerated. Surely it is not true that sense experience can be trusted only when verification by all the other senses is possible. If that were so, we could not justifiably believe we see a beam of light that is perceptually accessible only to our vision.

### *Problems confronting the experientialist approach*

There are many other relevant questions. Take first a psychological one of the kind relevant to epistemology. Do people ever really believe directly that, say, God is speaking to them, or is such a belief based—even if not self-consciously—on believing that the voice in question has certain characteristics, where one takes these to indicate God's speaking? Second, how is the possibility of corroboration by others—what we might call social justification—relevant? Does it, for instance, matter crucially, for experiential justification for believing in God, that not just any normal person can be expected to see God in the beauty of nature, whereas any normal person can be expected to

see a green field? Or is this contrast blunted by the marked differences in perceptual acuity among clearly normal people, particularly in complicated matters such as aesthetic perception in music and painting, in which what is directly heard or seen can be seen or heard only with both practice and sensitivity?

A related question is the role of testimony as a social source of direct justification. If it is true that (as argued in Chapter 7) beliefs based on testimony are commonly direct, then perhaps certain theistic testimony by some people provides knowledge of God to others. Even if one supposed that very few have theistic knowledge or justified theistic belief (at least “first-hand”), one might argue that the relevant testimonial chains can extend to many people—either during a given period of time or, where there is a community of believers, across time extending as long as thousands of years.

To be sure, justification seems different from knowledge here, at least insofar as we must have justification for believing someone to acquire justification for what is attested. But perhaps religious believers often have this justification for accepting testimony in religious matters; it is at least not obvious, for instance, that in order to be justified in religious beliefs on the basis of testimony they must have a kind of justification that is out of their reach as rational persons.

Whatever the place of testimony in providing theistic knowledge or justification, one might expand the possibilities for direct experience of God. Might God be seen, not necessarily in the ethereally direct way mystics have sometimes imagined, but in a more ordinary, if no less direct, fashion? If so, there is more ground to testify from as well as less need for testimony as a source of theistic knowledge or justification. Might God be seen, for instance, in nature, rather than so to speak inferred from it? Here is one of Gerard Manley Hopkins’s poetic expressions of that idea:

The world is charged with the grandeur of God.  
It will flame out, like shining from shook foil;  
It gathers to a greatness, like the ooze of oil  
Crushed . . .  
And for all this, nature is never spent;  
There lives the dearest freshness deep down things . . .<sup>24</sup>

After all, if nature is God’s work—perhaps God’s ongoing work—might there be a sense in which God is seen in it by those with appropriate sensitivity? A special sensitivity is needed even for seeing the beauty in a painting. To be sure, the relation of beauty to a painting that has it is different from the relation of God to nature conceived as revealing God. The point, however, is only that special sensitivity may be required for theistic perception, not that it is exactly like aesthetic perception.

The suggestion is not that nature is partly constitutive of God, at least not in the way that the shape and texture by which I perceive a spruce tree are

in part constitutive of it. Still, could nature, as some views apparently have it, be in some way part of God? If it is (as it in some sense would be on the view that the physical universe is God's body), then directly perceiving God may in a way be too easy, or at least quite easy to do without directly perceiving the divinity in what one sees. One could not see a beautiful landscape without seeing God, though one could see it without seeing it as manifesting God.<sup>25</sup> As stressed in examining perception, seeing something that has a property does not entail seeing it as having it. In this, divinity is no different from beauty.

The dimensions of these questions quickly widen, and even the many points that have come to light do not enable us to determine with any confidence whether there can be directly justified religious beliefs. It has so often been taken to be obvious that there cannot be, however, that it is important to see why it is really not obvious. It is at best very difficult to establish absolute restrictions on what sorts of beliefs can be directly justified. This holds even if the only way in which beliefs can be directly justified is by virtue of their grounding in the basic sources of justification.

A parallel point holds for absolute restrictions on what we can justifiably believe (or know) on the basis of one or more arguments. It is particularly difficult to determine what can be justifiably believed (or known) through a combination of plausible but individually inconclusive arguments for the same conclusion. As both coherentists and moderate foundationalists are at pains to show, there are times when a belief is justified not by grounding in one or more conclusive arguments, but by its support from—which implies some degree of coherence with—many sets of independent premises none of which, alone, would suffice to justify it.<sup>26</sup> The arguments that may work together here are not limited to the traditional kind proceeding from premises about the external world. Where one has non-inferential justification, say from a perceptual experience, one may formulate an argument that proceeds from premises describing the occurrence and character of the experience. Such arguments from experience can be combined with the traditional kind.

It must be granted, however, that it is often hard in practice to distinguish, even in our own case, between beliefs that are grounded directly in one of the basic sources and beliefs that are grounded in those sources through other beliefs of which we may not even be aware, or through inferences we do not realize we are making from propositions which we are aware we believe. This means that what we take to be direct belief, such as a belief that God has called on one to make sacrifice for someone else, may really be based on at least one other belief and may depend for its justification on the evidence or grounds which some other belief expresses. Still, even if we cannot tell whether a belief is inferential, we may be able to determine what further beliefs it is based on if it is inferential, and we may then be able to defend its justification on the basis of those.

Suppose for the sake of argument that there cannot be directly justified religious beliefs of the kind we have been discussing. There might still be

direct knowledge of such propositions, if (as I have argued) one can know certain kinds of things by virtue of a connection with them even if one does not have justification for believing them. For some religious people, even knowledge without justification might be considered very precious in this case. It would, perhaps, be one kind of faith.

### *Justification and rationality, faith and reason*

Our topic in these passages about theistic justification is sometimes called the question of faith and reason. In discussing that question, reason—above all rationality in holding religious beliefs—is commonly thought to be roughly equivalent to justification. I take it, however, that although a justified belief must be rational, a rational belief, while it presumably cannot be patently unjustified, need not be unqualifiedly justified.<sup>27</sup> Consider a belief that someone likes you. It can be rational on the basis of a vague “intuitive” sense before it is justified by evidence.

Moreover, justification seems tied more to specific justifiers than rationality is to any analogue of a justifier. If I justifiedly believe there is a cold glass in my hand, my justification is (chiefly) my tactual sensations; if I rationally believe that a painting is beautiful, there need be nothing comparable in the way of a sensory ground. I must have color sensations, but there is no sensation specifically of beauty as there is of the cold glass.

Perhaps rationality belongs for the most part to beliefs (and other elements, such as actions) that are (roughly speaking) fully consistent with reason and so, for instance, not obviously false, not crazy, and not mere results of wishful thinking. Moreover, rationality is to be understood by contrast chiefly with beliefs (and other elements) that are *irrational*,<sup>28</sup> whereas what is justified contrasts chiefly with what is *unjustified*. An unjustified belief—as many philosophers have discovered from their critics—need not be irrational. It is far easier to avoid having irrational beliefs than to avoid having unjustified ones. Rationality in a belief is achievable with lesser grounding than is required for its justification.

There is, then, at least one respect in which justification represents a less permissive normative standard than rationality. Mere absence of conditions that would make a belief that *p* unjustified does not imply that it is justified, but at most that one may suspend judgment on not-*p*, as opposed to being justified in believing that proposition. But in a rational person, absence of conditions that would make it irrational to hold a belief does tend to imply that it is rational. Normally, I may rationally believe a painting is beautiful if it seems so to me and I can find no reason to the contrary; I cannot justifiedly believe this without some substantially supportive ground that goes beyond the kind of impression that may suffice for rationality.

The suggested view is close to what has been called *epistemic conservatism*. In one form this is the view that our beliefs are, as it were, innocent unless “proven” guilty; in another form, it is the view that if a proposition

seems to one to be true, one is justified in believing it. The first view is too strong, since it applies to just any belief one might hold. The second view, a *phenomenal conservatism*, is plausible provided the seeming true is to some degree evidential, as it may often be. But the view suggested here is weaker in applying only to rationality rather than justification, in which the former is a more permissive (less strong) normative standard.

If rationality is a weaker—i.e., more permissive—normative notion than justification, it still provides a significant positive status that a theistic belief can have even if it is not justified. This is an important point. Scientific, moral, and other kinds of beliefs may also achieve rationality more readily than justification, even if, when they do, this is commonly a stage on the way to justification.

In any case, if rationality is possible without justification yet is implied by it, a plausible conclusion is that the experiential and rational grounds that, when sufficiently weighty, produce justification may, even when not quite weighty enough to yield justification, still render a belief based on them rational. A theistic belief might then be rational even if not justified. There might, to be sure, have to be some consideration weighing in the direction of justification, and one could speak here of some degree of justification; but as examples we explored earlier show, one can have some degree of justification for a proposition without having overall justification for believing it.

These points about the difference between justification and rationality do not show that anyone does hold rational theistic beliefs, or even that scientific or moral beliefs are ever rationally held. But if rationality is a weaker notion than justification, there would at least be better reason to think that this is so than there would be if the requirements for rationality were as strong as those for justification. In particular, whatever the weight of the considerations favoring the possibility of justified scientific, moral, and theistic beliefs—and I think the weight is substantial—those considerations weigh more heavily in favor of the possibility of rational scientific, moral, and theistic beliefs.

### *Acceptance, presumption, and faith*

One further line of thinking should be introduced here. We need not explore either justification or rationality in these three domains only in terms of belief. Belief has been utterly dominant in most epistemological discussions of cognition, but it is not the only cognitive attitude that raises epistemological questions or is appraisable in relation to justification or supporting grounds. There are attitudes weaker than belief in the degree of conviction they imply, yet strong enough in that psychological dimension to guide thought and action. Some philosophers have taken acceptance in this way. Accepting a scientific hypothesis, in this terminology, does not imply believing it, but it can commit one to using the hypothesis—say, that a certain illness is caused by a particular chemical—as a premise in (tentative) reasoning and in guiding one's day-to-day actions.

Similarly, in ethics one might presume the truth of a moral proposition, say that a certain job would involve one in a conflict of interest, without believing it. And in theology, one might have faith that, for instance, God is sovereign in the universe, without unqualifiedly believing this—though of course one may not have faith that this is so if one disbelieves it. In each case—acceptance, presumption, and faith—one cannot simultaneously have extremely strong doubts about the proposition; but one can have or entertain some degree of doubt, in a way one cannot if one genuinely believes it.<sup>29</sup>

It seems clear that the weight of evidence or grounding required for either justification or rationality will be less for these non-belief-implying attitudes than it is for belief. For instance, faith that a friend will recover from a disease can be rational when the situation is too bleak for justified (or even rational) belief that the recovery will occur. I might be perfectly reasonable, so far as the evidence goes, in having faith when I would be unreasonably underweighting the evidence if I believed the recovery will occur. And I can accept a hypothesis, at least for purposes of determining how to think and act in an urgent matter, when it would be premature to believe it. Acceptance and presumption can yield many of the practical benefits of believing but do not entail it. To be sure, religious faith differs in significant ways from the kind just described, but the main point still applies: whatever the grounds needed for justified theistic belief, weaker grounds will suffice for theistic faith with the same content, say that God is sovereign.

It turns out, then, that epistemology broadly conceived may consider not just the scope of our knowledge and justified belief but also the scope of our rational belief and even of other rational attitudes toward propositions, such as certain kinds of acceptance, presumption, and faith. This extension of epistemological appraisal to other, weaker attitudes provides more scope for rationality than there would be if belief were the only object of rationality. The same strength of evidence or grounding may take us further in the domain of attitudes like acceptance, presumption, and faith than in that of belief.

The question of how far our knowledge and justification extend beyond our beliefs grounded directly in experience or reason turns out to be complicated. We at least have warrant for rejecting the stereotypic view that whereas there obviously exists scientific knowledge as an upshot of proof, it is at best doubtful that there is any moral knowledge, or even can be religious knowledge. It seems a mistake to talk of scientific proof at all if that means (deductive) proof of scientific hypotheses or theories from premises indicating observational or other scientific evidence. Moreover, scientific knowledge does not often represent uncontroversial beliefs of precise generalizations, but is commonly either approximate knowledge, often known to need refinement, or knowledge of approximations, formulated with the appropriate restrictions left unspecified.

There is good reason to think that we also have, and certainly have not

been shown not to have, moral knowledge. And there is apparently no cogent reason to deny the possibility of religious knowledge. The same holds for moral and religious justification; and in all three instances, the scientific, the moral, and the religious, the case for the possibility of rational beliefs seems undefeated and, beyond that, stronger than the case for justification. Both cases appear still stronger as applied, not to beliefs, but to attitudes like acceptance, presumption, and faith, which in certain ways are weaker than belief. There are, of course, important skeptical arguments we have not considered, arguments that attempt to undermine all these positive conclusions and various other views about the scope of knowledge, justification, and rationality. It is time to examine some of those arguments.

## Notes

- 1 Cf. W.V. Quine's point that "Observation sentences [such as 'Rain' (for 'It is raining') and 'Milk' (for 'That is milk')] are . . . the vehicle of scientific evidence . . . But they are also the entering wedge in the learning of language." See *Pursuit of Truth*, p. 5.
- 2 The point that scientific theories are not proved, nor even, strictly speaking, disproved, by the kinds of observations that confirm or disconfirm them is developed by Pierre Duhem in *The Aim and Structure of Physical Theory*, trans. by Philip P. Wiener (Princeton, NJ: Princeton University Press, 1954); and the idea which explains this—that theories are tested in conjunction with a number of related hypotheses and assumptions rather than in isolation—is often called the *Duhem–Quine thesis* (or sometimes the *Quine–Duhem thesis*, given Quine's more detailed development of the view). For a short statement of the view see Quine, *Pursuit of Truth*, and Carl G. Hempel, *Philosophy of Natural Science* (Englewood Cliffs, NJ: Prentice-Hall, 1966).
- 3 For a good non-technical discussion of scientific inquiry and confirmation, see Carl G. Hempel, *ibid.*
- 4 This case closely follows one described on National Public Radio on November 25, 2009.
- 5 I assume here that one can *unknowingly discover* a truth, somewhat as one can unknowingly discover a diamond, thinking it may be paste (a zirconium imitation). Discovering that *p* does not entail knowing it and is compatible with thinking it is just a hypothesis to be tested (though not with believing not-*p*). One could presumably make a great discovery that is recognized *as* a discovery only by others.
- 6 My characterization of inference to the best explanation quite properly does not require that the inferred hypothesis be true. For such inferences are meant to be confirmatory and are not normally taken to entail truth. If, however, we make the plausible assumption that only true propositions are (genuine) explanations, then the term 'inference to the best



explanation' can mislead: the idea is roughly that of inference to what *seems* the best explanation (usually with the presupposition that the explanation is not just possibly correct but is adequately explanatory if true).

- 7 Feminist epistemologists have emphasized the social and cultural "situatedness" of knowers and their interdependence. They have also challenged the claims to scientific objectivity that many philosophers have made for scientific work properly carried out. For wide-ranging essays in and about feminist epistemology, see Kathleen Lennon and Margaret Whitford (eds.), *Knowing the Difference: Feminist Perspectives in Epistemology* (London: Routledge, 1994); and *The Monist's* special issue, *Feminist Epistemology: For and Against*, 77, 4 (1994).
- 8 See Helen E. Longino, 'The Fate of Knowledge in Social Theories of Knowledge', in Frederick Schmitt (ed.), *Socializing Epistemology: The Social Dimensions of Knowledge* (Lanham, MD: Rowman and Littlefield, 1994), reprinted in Jack S. Crumley II (ed.), *Readings in Epistemology* (Mountainview, CA: Mayfield, 1999), pp. 158, 159. For more recent wide-ranging works that treat feminist epistemology, see Miranda Fricker, *Epistemic Injustice: Power and the Ethics of Knowing* (Oxford: Oxford University Press, 2007), and Lorraine Code, 'Feminist Epistemology', in Jonathan Dancy, Ernest Sosa, and Matthias Steup (eds.), *The Blackwell Companion to Epistemology*, 2nd edn. (Oxford: Blackwell, 2010).
- 9 See Elizabeth Anderson, 'Feminist Epistemology: An Interpretation and a Defense', *Hypatia* 10 (1995), 50–84, reprinted in Crumley, *Readings in Epistemology*, p. 505.
- 10 *Ibid.*, p. 505.
- 11 For a representative collection of essays in this field, see Schmitt (ed.), *Socializing Epistemology*. The field is interdisciplinary, particularly in raising issues for philosophers, sociologists, psychologists, and others. For a wide-ranging study see Steve Fuller, 'Recent Work in Social Epistemology', *American Philosophical Quarterly* 33, 2 (1996), 149–66.
- 12 The paper by Hardwig cited in Chapter 7 notes an example of a scientific article with scores of co-authors, each of whom may be presumed essential for some item of knowledge important to the overall knowledge the article presents.
- 13 The wrongness is *prima facie*, not absolute, but this need not be made explicit, as it is ordinarily presupposed—and is certainly presupposed here.
- 14 That there is such moral knowledge is a major thrust of what is sometimes called 'Cornell Realism', since so many philosophers at or associated with Cornell University have defended it. For some of the seminal papers and other relevant studies see Geoffrey Sayre-McCord (ed.), *Essays in Moral Realism* (Ithaca, NY: Cornell University Press, 1988); and for a wider treatment of these issues see Walter Sinnott-Armstrong and Mark C.



- Timmons (eds.), *Moral Knowledge?* (Oxford: Oxford University Press, 1996).
- 15 This is the kind of thing W.D. Ross and other intuitionists have said about basic moral principles: they are intuitively knowable and self-evident, though seeing their truth may take a good deal of reflection. See, for example, Ross's *The Right and the Good*, esp. chapter 2.
  - 16 See, e.g., George Bealer, 'Intuition and the Autonomy of Philosophy', in Michael DePaul and William Ramsey (eds.), *Rethinking Intuition* (Lanham, MD: Rowman and Littlefield, 1998); Ernest Sosa, 'Minimal Intuitionism', in DePaul and Ramsey, *Rethinking Intuition*; and Michael Huemer, *Ethical Intuitionism* (Houndmills, Basingstoke: Palgrave-Macmillan, 2005). Sosa's notion may allow that both beliefs and seemings can count as intuitions (see, e.g., pp. 258–9); but for Huemer intuitions are a distinct, non-doxastic type of propositional attitude, evidenced when someone admits that *p* is intuitive (e.g., seems to be true), but denies believing it.
  - 17 See esp. Immanuel Kant, *Foundations of the Metaphysics of Morals*, trans. by Lewis White Beck (New York: Liberal Arts Press, 1959), section 2. Kant offered other formulations of the imperative (and the one I give involves some interpretation, e.g. the insertion of 'rationally'), but nothing said here turns on the interpretation).
  - 18 The wrongness in question is *prima facie*: this allows that excuses are possible but is compatible with the prohibition's being very strong, so that, for example, inexcusable violation is a grave moral offense.
  - 19 In chapter 2 of *Utilitarianism* (London, 1861; Indianapolis: Hackett, 2001), shortly after introducing his utilitarian principle, Mill contends that some pleasures are preferable to others, offers an empirical way to determine which of two pleasures is better (paragraph 5), and attempts to provide a similar way to weight quality of pleasure against quantity (paragraph 8). All this complicates his moral epistemology but does not alter the basic features in question here. A further complication is the possibility that he is best read as holding that the maximization of happiness standard applies to rules rather than individual acts. No major epistemological point made here turns on these issues of interpretation or on the indefiniteness of 'the relevant population'.
  - 20 Mill's attempted proof that pleasure is good is given in chapter 4 of his *Utilitarianism*; for an account of the view that goodness is to be determined by seeing what one would desire given adequate information and suitable reflection, see Richard B. Brandt, *A Theory of the Good and the Right* (Oxford: Oxford University Press, 1979).
  - 21 For an assessment of some of the most important arguments and an indication of the relevant literature see Keith Yandell, *Philosophy of Religion: A Contemporary Introduction to the Philosophy of Religion* (London: Routledge, 1999).
  - 22 For detailed discussion of the *sensus divinitatis* and an epistemology that

- gives it a substantial role, see Alvin Plantinga, *Warranted Christian Belief* (Oxford: Oxford University Press, 2000).
- 23 A more modest form of evidentialism, holding only that there is in fact no *actual* direct knowledge of God, would be unaffected by the bare possibility of the existence of an omnipotent God, but that evidentialist view is not our focus here.
  - 24 From Gerard Manley Hopkins's 'God's Grandeur', in A.J.M. Smith (ed.), *Seven Centuries of Verse*, 2nd edn (New York: Scribner's, 1957), p. 529.
  - 25 For an extensive treatment of the possible ways in which God may be perceived, with much general epistemology as context for demonstrating this possibility, see William P. Alston, *Perceiving God: The Epistemology of Religious Experience* (Ithaca, NY: Cornell University Press, 1991).
  - 26 It should be noted that noncognitivism can be applied to the philosophy of religion as to ethics. For many contemporary theological views, religious language is construed noncognitively, say as expressing reverent attitudes about people and their world, rather than as asserting metaphysical propositions about the ultimate origin and nature of the universe.
  - 27 I have compared and contrasted justification and rationality in 'Rationality and the Practice Conception of Justification', in Thomas D. Senor (ed.), *The Rationality of Belief and the Plurality of Faith* (Ithaca, NY: Cornell University Press, 1995).
  - 28 Bernard Gert has defended in detail the wider view that rationality is simply the contradictory of irrationality, which he takes as the prior notion. He might not accept, however, the use I make of a similar view here. See esp. *Morality* (Oxford: Oxford University Press, 1988).
  - 29 A detailed discussion of one of these non-belief cognitive attitudes is provided in my 'The Dimensions of Faith and the Demands of Reason', in Eleonore Stump (ed.), *Reasoned Faith* (Ithaca, NY: Cornell University Press, 1993).



# 13 Skepticism I

The quest for certainty

- **The possibility of pervasive error**

*Perfectly realistic hallucination*

*Two competing epistemic ideals: believing truth and avoiding falsehood*

*Some dimensions and varieties of skepticism*

- **Skepticism generalized**

*Skepticism about direct knowledge and justification*

*Inferential knowledge and justification: the problem of induction*

*The problem of other minds*

- **The egocentric predicament**

- **Fallibility**

*Three kinds of infallibility*

*Knowledge and fallibility*

- **Uncertainty**

*Knowing, knowing for certain, and telling for certain*

*Entailment as a requirement for inferential justification*

*Knowing and showing*

## 13 Skepticism I

### The quest for certainty

The commonsense view is that we all know many things. I believe I know many facts about my immediate surroundings, much about myself, something about the past, and a little about the future. I believe that we also have scientific knowledge, that we know some general moral truths, and that it is certainly possible that many of us know some religious truths. But there are reasons to doubt much of this. There are reasons to think that at best we know very little, perhaps just self-evident truths, for instance that if no vixens are males then no males are vixens, and a few propositions about our present consciousness, say that I am now thinking about the scope of human knowledge.

#### **The possibility of pervasive error**

As I consider these matters, I look back at the green field. I reassure myself that I see it vividly. I cannot help believing I do. But an inescapable belief need not be knowledge, or even justified. Suppose I am hallucinating. Then I would not know (through vision, at least) that the field is there.

#### *Perfectly realistic hallucination*

I find it impossible to believe that I am hallucinating. But I might find that impossible even if I were, provided the hallucination was as vivid and steady as my present visual experience. I begin to wonder, then, whether I really know that I am not hallucinating. If I do not know this, then even if I am in fact not hallucinating, can I know that there is a green field before me? Similarly, if I do not know that I am not simply having a vivid dream in which it seems to me that there is a green field before me, can I know that there is one there?<sup>1</sup>

Remembering that we can justifiedly believe something even if we do not know it, I think that at least I may justifiedly believe that there is a green field before me, even if I do not know that I am not hallucinating one (or merely “seeing” one in a dream). Moreover, if I justifiedly believe that there is a green field before me, how much does it matter whether I also know this? As we

saw in considering the value of knowledge and of this kind of justified true belief, the latter has substantial inherent value, even if (other things being equal) less than knowledge. Moreover, the likelihood that my belief is true, so far as that likelihood is something I can discern, depends on how probable the presence of the field is, given the sensory experience on which my belief is based; and in my attentiveness and caution as an observer, I have contributed all I can to that probability. Despite the possibility of hallucination, then, it appears that my belief remains justified, and it is as likely to be true as I can make it by any steps in my power, such as more carefully observing the grassy texture. Internally, in my own consciousness, I am being perfectly reasonable in continuing to believe that there is a green field there. So far as justification is concerned, I am beyond reproach.

These points about justification are plausible, but they give false comfort. Doubtless, we can have beliefs which, though they do not constitute knowledge, are justified, and we can have such a belief even if its basis is hallucinatory. But it is now not merely possible that I am hallucinating: I am also quite aware that I could be. Given this awareness, am I still justified in believing that there is a green field there? Should I not regard this belief as unjustified, suspend judgment on whether the field is there, and merely hope that it is?

### *Two competing epistemic ideals: believing truth and avoiding falsehood*

These questions produce a tension. I want to believe that the field is there if it truly is, for I have a deep-seated desire to believe as many significant truths as I can. But I also want to avoid believing that it is there if it is not, for I have a deep-seated desire to avoid believing falsehoods. Both of these desires are important; and they represent ideals that govern much of our thinking. But the two ideals pull against each other. The former inclines us to believe readily, since we may otherwise miss believing a truth; the latter inclines us to suspend judgment, lest we err by believing a falsehood.

The former ideal, calling on us to believe truths, pushes us toward credulity: believing on grounds that evidentially are too thin—or without grounds at all—and thereby believing too much. The latter ideal, calling on us to avoid believing falsehoods, pushes us toward a kind of skepticism: believing only on conclusive grounds, and thereby—if common sense is right about the matter—believing too little.

How can we balance these ideals with each other? So far, I have spoken more about how we fulfill the ideal of believing as many significant truths as we can than about how we might fulfill the ideal of avoiding belief of falsehoods. Clearly, the easiest way to fulfill the latter would be to suspend judgment on every proposition one entertains, or at least on those which, unlike certain self-evident truths, lack a luminous certainty that tends to compel assent. This is the kind of response characteristic of Pyrrhonian skepticism, an ancient variety tracing to Pyrrho of Elis (c. 360–275 BC).<sup>2</sup>

These reflections about possible error through hallucination, about the apparent vulnerability of justification in the face of such possibilities, and about the ideal of avoiding error suggest why philosophers have been so concerned with skepticism. In very broad terms, skepticism is most commonly conceived by philosophers roughly as the view that there is little if any knowledge. Call this view *knowledge skepticism*.

A related kind of skepticism is constituted by an attitude or feature of temperament, such as a disapproval of believing without conclusive grounds. This is not our direct concern. But if philosophical skepticism is not justified, then some common skeptical attitudes are not either, and some people who go through life with a skeptical attitude lack the kind of intellectual balance that goes with epistemic virtue. One reason, then, for studying skepticism is to approach a mean between two cognitive traits—intellectual vices, in the language of virtue epistemology. One vice is (excessive) credulity, which is too weak a disposition to doubt or withhold belief; the other is (excessive) skepticism, which is too strong a disposition to doubt or to withhold belief.

Skepticism may also target justification. Typically, skeptics do not take our justified beliefs to be significantly more numerous than our beliefs constituting knowledge. Call the view that we have little if any justification for belief *justification skepticism*. How far-reaching might a plausible skepticism of either kind be, and how is skepticism to be assessed? I want to pursue these questions in that order and at some length.

It may seem that skepticism offends so blatantly against common sense, and so lopsidedly prefers the ideal of avoiding falsehood over that of believing truths, that it should be dismissed as ridiculous. But it will soon be evident that skepticism is a serious, perhaps even irrefutable, challenge to common sense. Moreover, even if skepticism turns out, as phenomenalism apparently does, to be quite implausible, we learn a great deal about knowledge and justification from studying it.

A serious exploration of skepticism, whether or not we finally accept some form of it, also tends to help us to avoid dogmatism about our own personal views and a self-satisfied assurance that our collective outlook as rational observers of the world embodies knowledge of the sorts of things we think it does: facts about ourselves, our surroundings, and the ways of nature.

### *Some dimensions and varieties of skepticism*

To understand a skeptical view we should locate it in relation to at least four dimensions: (1) subject matter, say the past or the future or physical objects or other minds; (2) epistemic attitude, such as knowledge, justified belief, and suspended judgment; (3) modality, above all contingency or necessity, or the empirical versus the a priori; and (4) the kind of being it purports to limit, say human, subhuman, or superhuman. Regarding subject matter, my concern is wide-ranging. As for (2)–(4), my concern is with human beings and mainly with knowledge and justification regarding contingent empirical propositions.

Much skepticism, whether about knowledge or about justification, is restricted to a given kind of subject, for instance to propositions about the world outside oneself, or about the past, or about the future, or about ethics, religion, or science. Skeptical views also differ markedly in the status of the knowledge, and in the degree of the justification, they concern. A strong skepticism regarding propositions about the past, for instance, might hold that there is no knowledge, or even justified belief, about the past. A weaker skepticism might hold that although some beliefs about the past are justified to some degree, there is neither certain knowledge of the past nor any beliefs about it that are sufficiently justified to make it more reasonable to hold them than to suspend judgment on them.

Still another difference between skeptical views concerns their order. The usual skepticism is first order: it concerns the sorts of beliefs or knowledge we have discussed as typical of the kinds grounded in experience or reason, and not beliefs or knowledge about such beliefs or knowledge, say beliefs that ordinary perceptual beliefs often do constitute knowledge. First-order skepticism might deny, then, that I know there is a cold glass in my hand, even when I have the seemingly familiar experience I would describe as smelling the mint in my iced tea and feeling the cold glass in my hand. Second-order skepticism might say that even if I do know this, I do not know that I know it.

A first-order skeptic is committed to second-order skepticism: to holding, for instance, that there is no second-order knowledge to the effect that there is (first-order) knowledge, say knowledge of people, places, or things. This second-order skepticism is obviously true if there is in fact no first-order knowledge—as from that it would follow then that no one knows there is. But a second-order skeptic can also hold that even if there is first-order knowledge, no one knows this.

It is, moreover, natural for skeptics to hold their main views as necessary truths, as, for one thing, they commonly believe that for fallible creatures like us there cannot be knowledge or justification of certain kinds. I do not intend to discuss skepticism in detail in each of the many forms described, but what follows will apply to a very wide range of cases.

## **Skepticism generalized**

The skeptical challenges I have brought forward can be directed against all our beliefs about the external world, all our memory beliefs, all our beliefs about the future, and indeed all our beliefs about any subject provided they depend on our memory for their justification or for their status as knowledge. Memory is, after all, at least as liable to error as vision.

### *Skepticism about direct knowledge and justification*

Plainly, if all of the senses can deceive through hallucination, then beliefs grounded in any of the senses may be justificational or epistemically



undermined in the same way my belief that there is a green field before me may be undermined by a realization that I might have been hallucinating. Quite apart from whether perceptual beliefs are true, skeptics tend to claim that the possibility of such hallucinations either prevents these beliefs from being justified or, even if they are justified, precludes their constituting knowledge.

Suppose, for instance, that I might be having an auditory hallucination of bird songs. Then my present experience of (apparently) hearing them might not justify my believing that there are birds nearby and is certainly not a sufficient basis for my knowing there are, even if it is true that there are. Similarly, there is a counterpart of hallucination for memory beliefs: memorial hallucination, we might call it. I may have the memorial impression that when I was four I saw my parents kissing under the mistletoe, but this could be just a romantic fantasy masquerading as a memory.

Beliefs about the future are rather different from memory beliefs. The former concern future events and hence are not grounded in experiential states that in some way causally derive from things about which we have knowledge (as with perception). But even if there is no counterpart of memorial hallucination, there are equally undermining possibilities. For instance, a confident belief that I will talk with Jane could be a product of wishful thinking, even when in fact it is grounded in my long-standing intention to talk with her. Perhaps the belief is an anticipatory delusion. Even my belief that I will live to discuss skepticism could be mistaken for many sorts of reasons, including dangers to me of which I am now unaware.

Now consider our apparent general knowledge, whether a priori or scientific, say in arithmetic or science. Because it is possible to misremember propositions, or to seem to remember them when one does not, or to have a kind of memorial hallucination that gives rise to a completely groundless belief, it would seem that the only secure beliefs of general propositions are of the relatively few that we can know directly without needing evidence. This apparently leaves none of our general scientific beliefs, and only our a priori knowledge of self-evident propositions, epistemically unscathed.

### *Inferential knowledge and justification: the problem of induction*

Even if we leave aside problems about perceptual and memory beliefs, there is a difficulty for the commonsense view that justification or knowledge grounded (directly or indirectly) in a basic source can be transmitted inductively. The classical statement of this problem of induction—the problem of how to justify such inductive inferences—comes from David Hume.<sup>3</sup> Hume showed that one cannot know a priori that if the premises of a specific piece of inductive reasoning are true then its conclusion is also true. He noted that there is no contradiction in affirming the former and denying the latter. Moreover, one can conceive the premises being true while the conclusion is

false, whereas one cannot conceive its being true that (1) all human beings are mortal and Socrates is one of them, and yet false that (2) Socrates is mortal. Even good inductive reasoning is (deductively) invalid.

Consider the inductive reasoning from the premise that the sun has always risen each day to the conclusion that it will rise tomorrow. Of all such reasoning—reasoning “concerning matter of fact and existence”—Hume says:

That there are no demonstrative [roughly, valid, evidentially conclusive] arguments in the case seems evident, since it implies no contradiction that the course of nature may change and that an object, seemingly like those which we have experienced, may be attended with different or contrary effects.

(*Enquiry*, Sec. 4, Part II)

Hence, even if I do know that the sun has risen every day since time immemorial, I could be mistaken in believing that it will rise tomorrow, and it seems questionable whether I am even justified in believing this.

More generally, Hume’s arguments lead us to ask whether, if our premises could be true yet our conclusion false, we have any reason at all, on the basis of the premises, for believing the conclusion. And how can we ever know the conclusion on the basis of such premises? Indeed, how can we even be minimally justified in believing the conclusion on the basis of such premises? The problem of induction, as most often understood, is largely the difficulty of adequately answering these questions.

The problem can also be put in terms of probability. We normally operate on the commonsense presumption that when one thing is associated with another, say a sunrise with the passage of twenty-four hours, and the two have never failed to be associated in the same way, then the greater the number of cases of association, the greater the probability that the association will occur in a new case—for instance that the sun will rise tomorrow.<sup>4</sup> We also operate on the related commonsense presumption that for natural phenomena such an association can occur sufficiently often to yield justification for believing, and even knowledge, that the association will occur in a new case.

From a Humean perspective, it will not do to argue as follows: I am justified in believing my conclusion on the basis of inductive support for it, such as the past regular behavior of the sun, since past experience has shown that reasoning like this, which has had true premises, has also had true conclusions. For this way of defending an inductively based conclusion simply relies on yet another inductive argument—it gives a kind of inductive reasoning to support the view that certain kinds of inductive arguments justify one in believing their conclusions. It just inductively generalizes about inductive arguments themselves, using as a guide past experience in which we seem to have found out that by and large their conclusions turned out true when their premises were true.

That reasoning, then, apparently begs the question against Hume.<sup>5</sup> For

it assumes, without independent evidence, part of what he regards as false, namely, that inductive inference constitutes reasoning that either can ground knowledge of its conclusion, or can at least justify its conclusion, in the sense of providing good reason for it. We have taken the battle to a different field—that of inductive argumentation rather than sunrises—but we have added no new weapons nor enhanced our forces.

### *The problem of other minds*

One of the major points that Hume so powerfully defended—roughly, that non-deductive inferences are fallible—is by no means restricted to beliefs about the future. Such beliefs are, however, so prominent in his discussion of inductive inference that sometimes the problem of induction is narrowly conceived as that of how we can show that we have any reason to believe the future will be like the past. This conception is unduly narrow. Recall my observing Jim briskly shuffling papers and angrily mumbling curses. I cannot help believing, on this basis, that he is angry. But this reasoning leaves my belief clearly fallible: even if I know my premises (through perception), it does not follow that he is angry, and that could be false. He could be pretending.

The case of Jim's anger is alarmingly representative. Everything I believe about what is occurring in the inner lives of others seems to rest on grounds that are inductive in this way: what I observe—above all, their behavior—does not entail anything about their minds. They could be pretending, or psychologically abnormal; or some other source of error could occur. So if I cannot have knowledge of people's inner lives from their behavior, apparently I can never have it.

Worse still, if I cannot know anything about the inner lives of others, can I even know that there *are* others, as opposed to mere bodies controlled externally, or by hidden microscopic machinery, rather than directed through beliefs and intentions of the kind that I take to animate me?

There is, then, a problem of other minds. Can we know, or even justifiably believe, that there are any? If our experiences would be just as they are if the human bodies we interact with are controlled from outer space and have no inner life of their own, how can we know that those bodies are, as most of us cannot help thinking, animated by minds like ours?<sup>26</sup>

The problem is compounded when we realize that we can never directly verify, as we introspectively can in our own case, what is occurring in someone else's consciousness. Thus, all I can do to check on my inductively grounded beliefs about the inner lives of others is obtain further inductive evidence, for instance by observing whether they behave as one would expect if I am right in thinking them to be, say, angry. I cannot, as in my own case, introspectively focus on the events in their consciousness. How can I know anything about their mental and emotional life if I am in principle debarred from decisively verifying my beliefs about the contents and events of their consciousness? Even if I am sometimes right, I can never tell when.

It may be replied that by far my best explanation—and apparently a good explanation—of why other bodies behave as if they were animated by a mind is that they are so animated. The other hypotheses, such as control from outer space or by a machine, are far-fetched. The suggested reasoning sounds plausible, but notice that it is still a kind of induction: inference to the best explanation (abduction).

It is worth inquiring here whether some principle of inference to the best explanation is self-evident and may thus be presupposed in dealing with certain skeptical challenges. To simplify matters, let us bypass the common notion of the “best explanation.” Two explanations, after all, might be equally good. More important, our best explanation might not be good at all, hence of little help in supporting beliefs. Consider, then, this *moderate abductive principle*: if our only good explanation for a proposition we are amply justified in believing entails the truth or likely truth of a further proposition, we are *prima facie* justified in believing the latter proposition. Now imagine that (1) our only good explanation of why other bodies behave as if they were animated by a mind is that they are so animated. Taken together with the abductive principle, this entails that (2) we do have *prima facie* justification for believing that (say) the human creatures we interact with in normal ways have minds.

This is a plausible argument, but skeptics will reject it on at least two counts. They will question whether the abductive principle is self-evident or, perhaps, even true; and they will certainly challenge our presupposition that we are justified in holding (1).

I grant that, if the abductive principle is self-evident, it is not self-evident *that* it is. (Arguing that it is self-evident would be a difficult task I cannot undertake here.) As for (1), surely it too is plausible. Is there any good explanation for the apparently purposive, mentally guided behavior of other bodies that does *not* entail their being animated by minds? I doubt it. Granted, these bodies *could* be biological robots controlled from outer space, just as I could be hallucinating them in the first place. But is there any reason to think these possibilities more than bare logical possibilities? I see none.

It should be added here that the indicated inference to propositions about other minds as best explaining observed behavior is supported by a strong argument from analogy: again and again, when my body behaves in a certain way under certain conditions, I am in a certain mental state, say in pain when I am burned and cry out; so (surely) the same behavioral pattern in another body is accompanied by a similar mental state. Other bodies are so much like mine in physical structure and observed behavior that they are very probably animated by minds like mine.

There is no need to deny either that positing other minds is our only good explanation of what we seem to know about other human bodies or that the analogical argument just sketched is strong. Still, from one proposition's being our only good explanation of another (in the sense of ‘explanation’ relevant here) it does not follow that the first is true; and the analogies between my body and others at most render probable, rather than entailing, that some

other body is associated with mental states as mine is with my mental states. Recall that there are other possible explanations (such as the hypothesis of control of other bodies from outer space, or by a powerful and clever evil genius); these explanatory alternatives, if true, would leave my experience exactly as it is. For another thing, some of these alternative hypotheses can well explain the analogies that otherwise seem compelling.

Another way to see the power of these skeptical hypotheses is to note that our experience does not discriminate between the skeptical scenario and the commonsense one. In that scenario, our experience would be just what it is if we were steadfastly hallucinating the external world, including even the human bodies we seem to see. The same holds if we are not hallucinating but the human bodies are externally controlled. How, then, can our experience justify us in believing that there is an external world or that there are other minds?

Putting the problem somewhat differently, if our experience *underdetermines* the truth of propositions we commonly believe about the external world, roughly in the sense that it does not decisively indicate their truth as opposed to the truth of skeptical (or other) alternative hypotheses that can explain our experience, how can our experience justify our believing such commonsense propositions? If it cannot, and if, as Hume plausibly argued, we also cannot know that proposition, how can we be justified in believing anything at all about the external world?<sup>7</sup>

It is only a short step from this full-scale attack on inductive inference to a problem of the body. If, as a skeptic might well hold, our apparent knowledge of our own bodies is inductively grounded, being based on perceptions and bodily sensations somewhat as beliefs about external objects are, then can we know, or even justifiedly believe, that we have a body? Could we not be steadily hallucinating even our own flesh?

It might be replied that thoughts, including my reflections on skepticism, necessarily require an embodied thinker. But that point would only imply that I have some kind of body, not that I can know anything about it. The point is also far from self-evident and is indeed denied by philosophers in the powerful tradition of Descartes. They hold that we (persons) are essentially mental (or spiritual), hence non-physical, substances.

In any case, even if it should be true that thoughts can occur only in an embodied thinker, the only embodiment needed might be a brain. Hence, on the skeptical view imagined, the most we could know is that we are embodied in some way, say in a brain. Whether that brain is itself embodied, or ever interacts with anything else, would be beyond our knowledge. Why, then, could I not be alone in the world, or perhaps a “brain in a vat” kept alive in a nurturing liquid and subjected to hallucinations that falsely convey the impression of normal life?<sup>8</sup> Call this *the envatment problem*.

## The egocentric predicament

In this way, skepticism can drive us into an egocentric predicament: a position that makes it seem clear that all we can (empirically) know about the world, perhaps all we can justifiedly believe about it as well, concerns our own present experience. Perhaps, for all I know, I am a lone conscious ego vividly hallucinating a non-existent physical world. The kind of view in question—that only oneself exists—is called *solipsism*, and it serves as a limiting case to be avoided.

Most skeptics have tended to push no further, or at least not to express very much doubt about our capacity to know propositions of two specific kinds: those about what is currently going on in our minds and at least those a priori propositions that are luminously self-evident. But skeptics can push further. Descartes, in the first of his *Meditations*, raised the possibility that there was nothing of which he could (justifiedly) be certain. Recall introspectively grounded beliefs, such as that I am thinking about skepticism. It seems possible that this belief is mistaken. If that is possible, how can I know that I am thinking about skepticism? If I know, I cannot be wrong. But here error is possible. Perhaps I do not even have knowledge of my own conscious states.

To make this sort of argument work with beliefs of self-evident propositions we must, I think, strain. Descartes may perhaps be read as holding that God, being utterly omnipotent, could have falsified even propositions of the sort I am calling self-evident. But could an omnipotent being bring it about that while some dogs are pets, no pets are dogs? I see no reason to think so. As Thomas Aquinas and many other philosophers have maintained, omnipotence is simply not the power to “do” things that are absolutely impossible.<sup>9</sup> Power is exercised within the realm of the possible: impossible “deeds” are not candidates for any being to do.

If one accepts this point, one might argue that there is no act of bringing it about that while some dogs are pets, no pets are dogs. Calling this an act misuses the vocabulary of action. Hence, the impossibility that an omnipotent being can bring it about does not imply that there is any act which that being cannot perform. This point, in turn, deprives the skeptic of a way to argue that beliefs of necessary truths could be false.

This reasoning may not settle the matter, but it is sufficiently plausible to warrant leaving aside skepticism concerning beliefs of luminously self-evident propositions. These propositions seem not only incapable of falsehood, but, in some cases, also incapable of even being believed without justification, at least when carefully and comprehendingly considered. Leaving such skepticism aside takes little from the skeptic in any case. If these are the only knowable propositions, then we can know nothing about our world, not even about our innermost consciousness. We are at best in an egocentric predicament.

## Fallibility

In appraising skepticism, I want to formulate some of the main principles that underlie it in what seem its most plausible forms. If they can be shown to be unreasonable, then the skeptical threat to the commonsense view that we have a great deal of knowledge and justification can at least be blunted. In formulating and assessing these principles, we should distinguish skeptical threats to the generation of knowledge (or of justification) from skeptical threats to its transmission. It is natural to start with questions about generation. If no knowledge is generated, there is none to be transmitted.

### *Three kinds of infallibility*

Is there really any reason to doubt that, normally, introspectively grounded beliefs constitute knowledge? It may be true that such beliefs could be mistaken, but what is a skeptic entitled to make of this? The skeptical argument that comes to mind here is based on what I will call *the infallibility claim about knowledge*: if you know, you cannot be wrong. If we simply add the premise that you can be wrong in holding a given introspective belief, say that you are thinking about skepticism, it would seem to follow that such beliefs do not represent knowledge. This kind of argument from fallibility, as we might call it, can be applied to just about every sort of proposition we tend to think we know.

If, however, we look closely, we find that the infallibility claim is multiply ambiguous. There are at least three quite different things it might mean, and hence really three different infallibility principles.

The claim, 'If you know, you can't be wrong', might have the meaning of:

- 1 It must be the case that if you know that something is true, then it is true (i.e., you cannot know something that is false).

Call (1) the *verity principle*, since it says simply that knowledge must be of truths (verities). Knowledge can never have a falsehood as its object. The claim might, on the other hand, have the meaning of:

- 2 If you know that something is true, then it must be true, that is, the proposition you know is necessarily true (i.e., you can know only necessary truths).

Call (2) the *necessity principle*, since it says simply that knowledge is of necessary truths. Knowledge never has among its objects any propositions that could possibly fail to hold.<sup>10</sup> The claim 'If you know, you can't be wrong' might also have the meaning of:

- 3 If you know that something is true, then your belief of it must be true, in the sense that your believing it (the fact that you believe it) entails

or guarantees its truth (i.e., only beliefs that cannot be false constitute knowledge).

Call (3) the *infallibility principle proper*, since in saying that only infallible beliefs constitute knowledge it connects with skepticism more closely than (1) or (2). Knowledge, it says, is never constituted by fallible beliefs, those that can have falsehoods among their objects.

Unlike (2), (3) implies nothing about the propositional or other objects of knowledge; instead, it restricts the kind of belief that can constitute knowledge. And by contrast with (2), (3) also allows for knowledge of contingent (non-necessary) truths, such as that I exist. This proposition can be false (that I exist is not a necessary truth); but my belief of it is infallible and therefore cannot be false. If I now believe that I exist, then it follows that I do now exist.

### *Knowledge and fallibility*

We can now assess the skeptical reasoning that employs the infallibility claim in one or another interpretation. I will be quite brief in discussing the first two; the third is the most controversial and most important for skepticism.

The verity principle, (1), is plainly true: one cannot know something that is false. In this sense, knowledge is infallible. If it is false that the maple is taller than the spruce, then I do not know it is. But if this is all the infallibility claim comes to, it provides no reason to conclude that I do not know that I am thinking (or that anything else I believe is not genuine knowledge). Granted, it must be true that if I know I am thinking, then I am. But that tells us nothing about whether I do know I am. The verity principle is itself a verity, but it does not advance the skeptical cause.

The necessity principle, on the other hand, principle (2), seems mistaken. Surely I know some propositions that are not necessarily true, such as that I exist (as noted earlier, it is not a necessary truth that I exist, as it is that vixens are female). Even skeptics would grant that I cannot falsely believe this, since my believing it self-evidently entails that I exist (non-existent things cannot have beliefs at all). It may indeed be impossible for me even to be unjustified in believing the proposition that I exist when I comprehendingly consider it, in which case I have in mind some sense of myself. (Descartes seems to maintain in *Meditation II* that this case is impossible.) The same holds, of course, for you in relation to your belief that you exist.<sup>11</sup>

It might seem that we may grant the skeptic that the only kinds of proposition that cannot be falsely believed are either necessary or the rare kind that cannot be unjustifiably believed when comprehendingly considered. But that would also be a mistake: any proposition entailed by there being at least one belief is incapable of being falsely believed. Anyone's believing this one—that there is at least one belief—would entail that it is true (though it might be possible, given certain logical deficiencies, to believe such a proposition without having justification for it, as is certainly possible for necessary truths of



mathematics). Such cases suggest that there is no simple relationship between the kind of proposition believable with infallibility and the conditions for knowing or justifiedly believing it.

Even if the necessity principle were true, however, a skeptic could not reasonably use it, without first defending it by adequate argument, against the commonsense view that introspective or even perceptual beliefs normally constitute knowledge. For clearly they are not beliefs of necessary truths, and defenders of common sense do not take them to be. Hence, invoking the necessity principle against common sense, without first arguing for the principle, would be in effect a flat denial that such beliefs constitute knowledge. That would beg the question against the commonsense view.

Suppose, for instance, that a skeptic says that if you know, you cannot be wrong, in which this means (2), then notes that introspective and perceptual beliefs (which are of propositions that are not necessary) can be false, and concludes that such beliefs do not constitute knowledge. This would not be presenting a good reason to believe the conclusion, but just denying, disguisedly given (2), the commonsense view that we have introspective and perceptual knowledge. There may seem to be a good argument here, because it is so easy to take 'If you know, you can't be wrong' as asserting the verity principle. But that principle is acceptable to common sense, whereas the necessity principle is not. To argue for the latter by allowing the plausibility of the former to serve as support for it is to trade on an ambiguity. It masks poor reasoning—or the absence of any argument or support at all.

The infallibility principle proper, (3), in effect says that only infallible beliefs can be knowledge. Now as we have seen, some beliefs of contingent propositions are infallible. Consider my belief that I exist, and my more specific belief that I have a belief. Just as my believing I exist entails that I do exist, if I believe I have a belief, it follows that I have one: I have at least that very belief even if I have no others. Beliefs like these might be called *self-grounding*, since comprehendingly considering them constitutes a sufficient ground both for justifiably holding them and indeed for their truth.

The infallibility of these two contingently true, but self-grounding, beliefs shows that despite appearances, (3) is not equivalent to (2), since (3), the infallibility principle, but not (2), the necessity principle, allows knowledge of propositions that are not necessary (i.e., are contingent propositions). But why should we accept (3)? What reason can the skeptic give for it? Not that if you know, you cannot be wrong; for when we look closely, we find that when plausibly interpreted as meaning (1), that is no help to the skeptic, and when interpreted as (2) or (3) it just flatly asserts the skeptical position against common sense.

What makes the infallibility claim seem to give the skeptic an argument against common sense is the way skepticism can trade on the ambiguity of that formulation: one finds the argument from fallibility attractive because its main premise, conceived as equivalent to (1), is so plausible; yet the argument

succeeds against common sense only if (2) or (3) is a legitimate premise. (2) is clearly mistaken, and it is doubtful that the skeptic has a cogent argument for (3). It will help to consider first the bearing of the concept of uncertainty, one closely related to the notion of infallibility.

## Uncertainty

Like fallibility, uncertainty has seemed to many skeptics to leave us with little, if any, knowledge. Recall the possibility that I am hallucinating a green field before me when there is none there. May I be certain, then, that there is one there? And can I ever tell for certain whether or not I am hallucinating? Skeptics tend to answer negatively and to contend that if we cannot tell for certain whether we are hallucinating, we do not know we are not hallucinating. They also tend to argue that if we do not know we are not hallucinating, surely we do not know that there is a field there.

Moreover, in a way uncertainty cuts deeper than fallibility: for even if I believe a theorem of logic that cannot be false and so have an infallible belief, I may not be justified in taking my proof to be sound and so cannot be justifiedly certain. Uncertainty arises when one's grounds are inconclusive, and it can arise, as with beliefs of theorems, even when one's belief is infallible. Thus, even infallibility is insufficient to render a belief knowledge. At least two important principles are suggested here.

One principle suggested by reflection on these questions about possible error is the *certainty principle*: *if one cannot tell for certain whether something is so, then one does not know it is so*. This principle is plausible in part because, typically, 'How can you tell?' and 'How can you be certain?' are appropriate challenges to a knowledge claim. Moreover, 'I know, but I am not certain' sounds self-defeating, in a way that might encourage a skeptic to consider it contradictory. Further support for the certainty principle can be derived from the idea that if our grounds for a belief underdetermine its truth—as when a skeptical possibility such as the Cartesian demon hypothesis can also explain our having those grounds—then one cannot tell for certain that the belief is true.

Another principle suggested by our questions about the possibility of hallucination is the *back-up principle*: a belief that  $p$  constitutes knowledge only if it is backed up by one's knowing, or at least being in a position to know, the falsity of any proposition inconsistent with  $p$ . Thus, if I believe that there is a field before me, then, as this proposition is inconsistent with my merely hallucinating a field, I know this proposition only if I am at least in a position to know that I am not hallucinating.

The back-up principle is plausible in part because one is in a sense responsible for the implications of what one claims to know. If, for instance, I claim to know that there is a green field before me, and that proposition implies that the field is not a pavement textured and painted to look just like a green

field, it would seem that I had better know or at least be in a position to know that it is not such a pavement. This, in turn, is commonly taken to imply that I must at least be justified in rejecting that strange possibility.

The upshot of this skeptical reasoning is that if I know that there is a green field before me, I apparently must be prepared to back that up by justifiedly rejecting exactly the sorts of possibilities that the skeptic reminds us are always there, in abundance. But must I be thus prepared? Let us consider the certainty and back-up principles in turn.

### *Knowing, knowing for certain, and telling for certain*

Chapter 10 argued that knowing does not imply knowing for certain. This conclusion suggests that the kind of certainty in question, epistemic certainty, is not required for knowledge, and that having such certainty may be something quite different from simply knowing. Still, from the point that knowing need not be knowing for certain, it does not follow that one can know without being *able* to tell for certain. Thus, the skeptic may still maintain that the certainty principle undermines the commonsense view that we have perceptual knowledge.

Let us first ask what it is to tell for certain. A skeptic may mean by this acquiring knowledge, in the form of an infallible belief, of a proposition that entails the truth of what one can tell is so. Thus, to tell (for certain) that one is not hallucinating a green field one might, like Descartes in the *Meditations*, prove that there is a God of such goodness and power that—since it would be evil for God to allow it—one could not be mistaken in a belief properly based on such a vivid and steadfast perception as one now has of a green field. We can tell for certain that there is an object before us because we can prove that God would not allow us to believe this under the present conditions unless it were true.

Some thinkers might embrace Descartes's theistic solution here. But one might also reject the skeptical principle in question, the infallibility principle. To require that a belief can be knowledge only if—whether in Descartes's way or a similar fashion—it can be conclusively shown to be true would again beg the question against the commonsense view that a belief can constitute knowledge without being infallible (a belief that can be absolutely conclusively shown to be true is infallible). Thus, if skeptics have no good argument for the principle of infallibility proper, they should not assume that principle in defending the view that we can know only what we can “tell for certain” in this strong sense of the phrase.

Perhaps, on the other hand, telling for certain is simply a matter of ascertaining the truth in question by some means that justifies one in being (psychologically) certain of what one can tell, even if not maximally certain (if there is a maximum here). If so, perhaps we normally can tell for certain that we are not hallucinating, for instance by seeing whether the senses of touch

and smell confirm our visual impression. To be sure, the confirming experiences do not entail that there is a green field before me. But we still have no good argument that certainty (or knowledge) may arise only from entailing grounds (another controversial view, shortly to be discussed). Thus, this point does not establish that confirming experiences cannot enable us to tell for certain that we are not hallucinating.

Moreover, suppose that we interpret telling for certain in the modest way just suggested, and that we can tell for certain in this sense that what we know is true. In that case, perhaps there is a weak sense in which beliefs constituting knowledge are infallible. They need not be such that it is absolutely impossible (logically impossible, in a broad sense) that they be false, as in the case of my belief that I exist. There need only be something about our grounds for them in virtue of which they (empirically) cannot be false, say because it would violate the laws of nature. Water cannot flow (as opposed to being pumped) uphill, but this is empirically impossible, not absolutely so, as it is impossible for some pets to be dogs without any dogs being pets. By contrast, it is not even empirically impossible to win a lottery with just one of a trillion tickets, and this can explain why the skeptically inclined will likely refuse to say one can know in advance that such a ticket will lose.

It may be true that grounds of what is commonly considered to be knowledge are typically such that, given those grounds, the belief constituting that knowledge cannot be false (at least cannot be false given the laws of nature). Suppose this is true. Should we now say to the skeptic that the beliefs commonsensically considered knowledge, such as many perceptual ones, are empirically certain? We may say this only if we keep in mind what was wrong with inferring the necessity principle from 'If you know, you can't be wrong'. There surely might be causal laws of nature which guarantee that *if* one is situated before a field in good light, as I am, and one has visual experiences like mine caused by the field as mine are, then one sees it, and hence cannot falsely believe that it is there. But this hypothetical proposition does not imply that my belief is, of empirical necessity, true, as a law of nature at least commonly is, any more than the "logical" law that it is necessary that if one knows that  $p$  then  $p$  is true implies that  $p$  itself is necessary. A guarantee of truth *given* certain grounds is not a guarantee of even empirically necessary truth, much less of epistemic certainty, any more than a guarantee of payment is a guarantee of payment in gold or in some medium that cannot be devalued.

If the existence of causal laws and associated causes of many of our beliefs implies the truth of many of our beliefs grounded in experience, such as my belief that there is a green field before me, it does not follow that those true ones cannot be mistaken, in the sense of being epistemically certain, or conclusively justified, or any other epistemic status high enough to satisfy a skeptic. What follows is only that given the laws of nature and their causal grounding, they are true. This seems more than enough for common sense.

The skeptic gives us no good argument to show that there are no such laws or that such law-based truth is not quite sufficient to render a belief knowledge whether it represents certainty or not.<sup>12</sup>

### *Entailment as a requirement for inferential justification*

The back-up principle fares no better than the infallibility principle proper. For one thing, it depends on the assumption, which defenders of common sense stoutly reject, that in order to know that something is true, one must have grounds that entail its truth. To see that the back-up principle depends on this, consider first a very simple case. Take the proposition that it is false that there is a green field before me. This is inconsistent with what I believe, namely, that there is one before me. Hence, the back-up principle requires that I at least be in a position to know that this is false. Its falsity entails that there *is* a green field before me.

The back-up principle may seem true because one may think: How else, besides being able to know the falsity of propositions incompatible with what I believe, can I be adequately armed against the threat of error? If I am not in a position to know that propositions plainly incompatible with what I believe are false, I cannot properly back up what I believe. But the falsity of the negative proposition that there is not a green field before me entails that there is one before me; for if it is false that it is false that there is one, then it is true that there is one. Thus, if, by virtue of how I must be able to back up my original claim, I do know that this negative proposition is false, then I thereby have (and know) an entailing ground for the truth of what I originally believed—that there is a green field before me.

Now take a case in which backing up what I think I know is more complicated. Consider the proposition that what I take to be a green field is really a pavement with such a realistic-seeming grassy green texture that I cannot tell (perceptually) that it is really not a field. Must I be in a position to know that this is false to know that there is a green field before me? The very description of the case suggests that I cannot know, at least by using the senses unaided by experimentation or specialized knowledge, that the field is not a textured pavement. But why must I be able to tell this at all? Is there any reason to think that the field might actually be dyed pavement? Is that a “relevant alternative,” some philosophers would ask?

One might object that in order to know a proposition I must be in a position to know whatever follows from it (or at least obviously follows from it). After all, if something does follow from what I know, I could infer it by valid steps from what I initially know, and thereby come to know it too.

This is an important objection. But in discussing the transmission of knowledge and justification, we considered cases that apparently undercut the objection. I can apparently know the sum of a column of figures even if I cannot, without further checking, know something which obviously follows from it: that if my wife (whom I justifiedly believe to be a better arithmetician) says this is not the sum, then she is wrong. If this can be true of me,

then neither knowledge nor justification is automatically transmitted across valid deductive inference.

It apparently will not do, then, to say that we can always count on the transmission of knowledge from propositions we know or justifiedly believe to those they entail, even when the entailment is, as in our example, obvious. Thus, even though my seeing a green field plainly entails (for instance) that I am not seeing a pavement with a textured grassy-looking surface, I presumably do not have to be in a position to know or justifiedly believe (by inferring it) that this proposition is false.

Suppose, however, that this view is mistaken, and that knowledge and justification are always transmitted across valid deductive inference. It may be plausibly argued that I do have justification for rejecting the skeptical hypothesis that there is a pavement before me textured to look just like a green field. It is not just that it appears to me that there is a green field before me; I also have no reason to think there is anything abnormal in the situation, and some reason to think that, in cases like this, large, nearby familiar kinds of things are as they appear to me in such vivid and careful observation. That they are as they appear is at least supported if the abductive principle is sound; for their being as they appear is surely my only good explanation of their appearing as they do. On balance, then, I may apparently reject the skeptical hypothesis and I do know or at least justifiedly believe that there is a green field before me.

We could also stress that the kinds of grounds I have for believing there is a green field before me are plainly sufficient for knowing this proposition and then take that very proposition as my premise for the entailed conclusion that there is not a pavement before me textured to look like grass. On this view, the point is that by virtue of perceptual justification we gain (commonsense) knowledge of a conclusive ground for rejecting the skeptical hypothesis.<sup>13</sup>

There are other factors one might cite in defending commonsense views of the scope of our knowledge, indeed, too many to discuss here. My point is simply this. Because the skeptic has not provided good reasons for the principles I have already rejected (or for comparably strong principles), even if knowledge and justification are always transmitted across valid inference, there may be good reason to say that skeptical hypotheses, such as that the "field" is a cleverly painted and textured pavement, may be justifiedly rejected.<sup>14</sup>

### *Knowing and showing*

There is something we may grant to the skeptic that will help to justify my rejection of the certainty and back-up principles. Admittedly, to show the skeptic that my original belief is knowledge, in the face of the suggestion that one of those explanations of its falsity holds, I may have to know that, and perhaps why, this explanation does not hold. Showing something, after all, commonly requires invoking premises for it, and one must presumably know or justifiedly believe those premises if one is to show a conclusion from

them.<sup>15</sup> The question 'Do you know?' tends to move discussion to a second-order context in which one seeks not simply to offer grounds for what one takes oneself to know, but grounds for the second-order proposition that one knows it. After all, a direct answer to, for instance, 'Do you know that she missed the train?' is something like, 'Yes, I know because I just checked the station', rather than 'I just checked the station'. The latter reply provides a ground on which one may know that she missed the train and only indirectly implies that I do know it. A skeptic would not grant this evidential power to such a ground, however, and would indeed not take my citing the station check to provide an adequate answer to 'Do you know that she missed the train?'

Still, we may ask, why, in the absence of the need to show that I know, must I, in order simply to have knowledge, have the capacity to show that I have it, as the back-up principle strongly suggests? Surely I need not. I can know that if some dogs are pets then some pets are dogs, even if I cannot show this self-evident truth—perhaps simply because I can think of nothing more obvious to use as a reasonable premise from which to show it. And if my wife raises no question of whether my arithmetic answer is correct, I can know that answer even if I cannot show—without obtaining further grounds for the answer—that I do know it. (If my original justification were good enough to enable me to show that if she says the sum is wrong then she is wrong, perhaps it would also enable me to know, even without showing it, that if she says this, she is wrong.)

The point that one can know without being able to show that one does drastically weakens the case for the back-up principle. Moreover, if, as seems quite possible, I can know the sum on the basis of my calculations without being able to show that I do—apart from gaining new evidence—then I can know it without being able, given my evidence from careful calculation, to tell for certain whether it is true. That would require new calculations and hence new evidence. This second point directly cuts against the certainty principle as well as against the back-up principle.

Examining the relation between knowing something and being able to show it also indicates that the converse of the certainty principle—the *show-know principle*, we might call it—should also be rejected: *being able to show something one believes, even being able to prove it, entails knowing it*. This can be seen from our example. Suppose I can now show that if she says the sum is wrong, she is mistaken, by doing a more careful calculation twice over. Does it follow that I now know this proposition for certain? I do not see that it does. From the fact that I now have the ability to show something I believe, it does not follow that I now know it at all. Having the raw materials to create something—here grounds for knowledge—does not entail already having it. Moreover, suppose that, as sometimes happens, I am lucky in a mathematical hunch; I might still be capable of constructing a proof I myself would not have expected to discover. By good fortune, we may have raw materials to create a foundation for something we have fabricated only by a stroke of luck.

My examples of showing something are cases of deductive demonstration.



But it would be a mistake to think that all instances of showing require deductive demonstration or proceed from absolutely conclusive grounds. Both points were discredited by our discussion of how propositions are established through scientific reasoning. It appears that we can know certain kinds of things even if we cannot show them and that, when we can show to be true something we know, the kind of showing possible for us need not require proof. We have yet to explore, however, whether there is any way in which we might show that a commonsense view of the scope of our knowledge is justified. This will be the main business of the next chapter.

## Notes

- 1 Some writers on skepticism prefer using the dream case rather than the hallucination one, perhaps in part because Descartes so famously used a dream argument in his *Meditations*. For relevant discussion, esp. of the dream argument, see, e.g., Barry Stroud, *The Significance of Philosophical Scepticism* (Oxford: Oxford University Press, 1984); Michael Williams, *Unnatural Doubts* (Oxford: Basil Blackwell, 1991); Robert Fogelin, *Pyrrhonian Reflections on Knowledge and Justification* (Oxford: Oxford University Press, 1994); Sanford Goldberg, *Anti-Individualism* (Cambridge: Cambridge University Press, 2008); and Ernest Sosa, *A Virtue Epistemology*, esp. chapter 1. Another way to raise skeptical possibilities is to imagine being a “brain in a vat,” i.e., that one’s brain is alive in a liquid and given just the sorts of stimulation it now has, so that one would seem to experience things just as one now does. This example derives from Hilary Putnam, who uses it to argue, against skepticism, that the very *content* of such sentences as ‘I am a brain in a vat’ prevents their being intelligibly thought in the way Descartes apparently believed possible. Because of how language and conceptualization work, “although the people in that possible world [in which they are merely brains in a vat] can think and ‘say’ any words we can think and say, they cannot (I claim) *refer* to what we can refer to. In particular they cannot think or say that they are brains in a vat (*even by thinking ‘we are brains in a vat’*).” See the selection from Hilary Putnam, *Reason, Truth and History* (Cambridge: Cambridge University Press, 1981), entitled ‘Brains in a Vat’, in Huemer, *Epistemology*. The literature contains much critical discussion, including Huemer’s ‘Direct Realism and the Brain-in-a-Vat Argument’, reprinted (from *Philosophy and Phenomenological Research* 61, 2 [2000]) in Huemer, *Epistemology*. See also Timothy Williamson, ‘On Being Justified in One’s Head’, and my response to him, ‘Internalism and Externalism in Epistemology and Semantics’, both in Mark Timmons, John Greco, and Alfred R. Mele (eds.), *Rationality and the Good: Critical Essays on the Ethics and Epistemology of Robert Audi* (Oxford: Oxford University Press, 2007).
- 2 Pyrrhonian skepticism need not imply that suspending judgment is



always psychologically possible; and it may also allow that one can accept a proposition for the sake of argument, and so, without believing it, act on it as one who believes it would act; but there is no need to discuss this position here. If what I say about skepticism in general is sound, it can be readily applied to the Pyrrhonian form.

- 3 See, for instance, section IV of Hume's *Inquiry Concerning Human Understanding* (first published in 1748).
- 4 This formulation is roughly the one given by Bertrand Russell in *The Problems of Philosophy* (Oxford: Oxford University Press, 1912). The formulation should be understood to apply only to phenomena of a certain kind, such as we explore in scientific inquiry and much everyday investigation; it will not hold for certain special cases. For instance, with an increase in the number of instances in which I lose a fair lottery in which I hold one of a million coupons, there is *no* change in the probability that I will lose; the probability remains the ratio of the number of tickets I hold—one—to the total number: 1 million. To think my good day is now more likely to come is to commit the *gambler's fallacy*.
- 5 I cannot take time here to consider begging the question in any detail; it is an important but elusive notion. For a detailed treatment see Walter Sinnott-Armstrong, 'Begging the Question', *Australasian Journal of Philosophy* 77, 2 (1999), 174–91.
- 6 A remark attributed to Bertrand Russell introduces (humorously) the possibility that we may, perhaps haphazardly, be sometimes right and sometimes wrong about other minds: "There may be other minds," he quipped, "but there certainly aren't many of them." In this spirit I might note another twist to the problem of other minds. As usually conceived, it concerns whether, in effect, there are as *many* minds as there seem to be. But reflection on skepticism can also lead us to a converse worry. How do I know that when I am asleep my body is not taken over by another mind, one connected, perhaps, with a part of the same brain as goes with my mind? And why might there not be several others who control this body when I do not? Granted there could be a conflict with another mind over, say, the movements of my right arm; but I could be built (or programmed) so as never to be conscious when another mind takes over this body. Call this the *problem of too many minds*.
- 7 There is much recent literature on the extent to which skeptical hypotheses undermine commonsense views about the extent of our knowledge. See, for instance, Butchvarov's *Skepticism and the External World* (New York: Oxford University Press, 1998), Moser's *Philosophy after Objectivity* (Oxford: Oxford University Press, 1993) and the books by Fogelin, Huemer, Klein, Nozick, Rescher, Sosa, Stroud, and Michael Williams also cited in the bibliography. For a helpful discussion of skepticism focusing on the underdetermination problem, see Jonathan Vogel, 'Dismissing Skeptical Possibilities', *Philosophical Studies* 70 (1993), 235–50.

- 8 This is the term widely used in recent literature in connection with a debate between, on the non-skeptical side, Hilary Putnam, who (as quoted in note 1) denies that one really could be both a brain in a vat *and* have the mastery of language needed to raise the question of skepticism, and, closer to the skeptical side in interpreting such examples, a number of other philosophers, including Anthony Brueckner, 'Trying to Get outside Your Own Skin', *Philosophical Topics* 23, 1 (1995), 79–111, which contains references to Putnam's original discussions of the brain-in-a-vat problem and a number of more recent discussions. See also the paper by Huemer cited in note 1 and his *Skepticism and the Veil of Perception*.
- 9 See, for example, Thomas Aquinas's *Summa Theologica* (written in the thirteenth century), Ia, question 25, a.3.
- 10 As an epistemic principle, (2)—and indeed (1) and (3) as well—would commonly be taken by proponents to be necessarily true. Thus (2) would rule out even the *possibility* of knowledge of falsehoods, as opposed to the mere occurrence of it. But for our purposes the formulations as more simply stated will serve.
- 11 That my thinking entails my existing does not, of course, entail that my existing entails my thinking. But Descartes's claim (also in *Meditation II*) that his essence *is* to be a thinking thing led to the following joke (which I recount as I remember it). *Bartender to customer*: Do you want another? *Customer*: I think not. *Outcome*: The customer disappears.
- 12 If there are such laws then there is empirical grounding that is conclusive in the sense that it implies the proposition it grounds with "natural necessity," the kind appropriate to causal laws. Because those are not absolutely necessary, as are logical laws and necessary truths as described in Chapter 5, it would not follow that the implication is an entailment.
- 13 This is the kind of strategy taken by Peter D. Klein in *Certainty: A Refutation of Scepticism* (Minneapolis: University of Minnesota Press, 1981); he extends it in 'Skepticism and Closure', cited in note 22 of Chapter 9 and reprinted in Huemer, *Epistemology*, pp. 552–74. In part, the issue concerns whether we may simply take it as obvious that we do know certain things before we have a *criterion* of knowledge, e.g. an account that tells us both what knowledge is and whether beliefs constituting it must be infallible. For detailed discussion of this *problem of the criterion*—the problem of whether cases of knowledge are prior to accounts or vice versa—see R.M. Chisholm, 'The Problem of the Criterion', in his *The Foundations of Knowing* (Minneapolis: University of Minnesota Press, 1982), reprinted in Huemer, *Epistemology*; and Robert Amico, *The Problem of the Criterion* (Lanham, MD: Rowman and Littlefield, 1994). For an examination of Moorean commonsensism in relation to criteria, dogmatism, and skepticism, see James Pryor, 'The Skeptic and the Dogmatist', *Noûs* 34 (2000), 517–49.
- 14 There is a sophisticated and plausible compromise with skepticism that deserves note here. One could argue that knowledge must be understood

not unqualifiedly but in terms of relativization to “contrast classes.” Thus, relative to the contrast class of near perfect imitations, I do not know that there is a green field before me; relative to the contrast class of ordinary green things encountered in the same visual way, such as green ponds and green canvas laid out for picnicking, I do. For a detailed statement of this view—which may be regarded as a kind of contextualism—see Walter Sinnott-Armstrong’s chapter on moral skepticism in Walter Sinnott-Armstrong and Mark Timmons (eds.), *Moral Knowledge?* (New York: Oxford University Press, 1996). One reply is that knowledge may be understood without relativization provided we at least (1) recognize that knowledge *attributions* occur in a context and are in some way relative to it, and (2) distinguish between the following kinds of things: (a) knowing there is a green field before one and (b) knowing that there is a green field before one *as opposed to* a pavement textured to look just like one. Knowing the former does not entail knowing the latter, contrastive proposition (nor does Sinnott-Armstrong claim this), but it is easily taken to require that because (i) the latter proposition apparently follows from the former and (ii) *noting* the possibility of ignorance of the latter is an intelligible challenge to a *claim* to know the former.

- 15 The ‘presumably’ is meant to allow that there is a sense of ‘show’ in which one need not have justification for one’s premises: if they are true and are cogent grounds for what one wants to show, then invoking them may serve to show it. Here, however, one might not be justified in believing the very thing one shows. This objective, externalist way of showing—call it *de facto showing*—something is not the one of interest here, which might be called *dialectical showing*, as it figures crucially in philosophical disputation.

## 14 Skepticism II

The defense of common sense in the face of fallibility

- **Negative versus positive defenses of common sense**
- **Deducibility, evidential transmission, and induction**  
*Epistemic and logical possibility*  
*Entailment, certainty, and fallibility*
- **The authority of knowledge and the cogency of its grounds**  
*Epistemic authority and cogent grounds*  
*Grounds of knowledge as conferring epistemic authority*  
*Exhibiting knowledge versus dogmatically claiming it*
- **Refutation and rebuttal**
- **Prospects for a positive defense of common sense**  
*A case for justified belief*  
*The regress of demonstration*  
*A case for knowledge*  
*A circularity problem*
- **The challenge of rational disagreement**  
*Intellectual pluralism*  
*Epistemic parity*  
*Dogmatism, fallibilism, and intellectual courage*
- **Skepticism and common sense**

## 14 Skepticism II

### The defense of common sense in the face of fallibility

The previous chapter indicates various ways in which a skeptical case against what might be called epistemological commonsensism can be resisted. But even if such resistance is warranted, it leaves open just what may be said positively in favor of the view that the kinds of apparent knowledge and justified belief discussed in the first 12 chapters of this book are genuine. This chapter will explore that question.

#### **Negative versus positive defenses of common sense**

In the context of thinking about skepticism, it is easy to forget that knowing something does not require being able to show that one knows it. For in thinking about skepticism we are likely to be trying to defend, against a skeptical onslaught, the commonsense view that there is much knowledge, and we easily think of defending this view as requiring us to show that there is knowledge. There is, however, more than one kind of defense. The two kinds I have in mind are analogous to standing firm as opposed to attacking.

A negative defense of common sense, one that seeks to show that skeptical arguments do not justify the skeptic's conclusion, does not require accomplishing the second-order task of showing that there is knowledge or justified belief. That achievement is required by a positive defense of common sense, one that seeks to show that we have the kinds of knowledge and justified beliefs common sense takes us to have. A negative defense requires only contending that skepticism provides no good argument against common sense.

It does not appear that skepticism as so far examined provides a good argument against common sense. Why, for instance, should the skeptic's merely suggesting a possible explanation of how there could be no green field before me, without giving any reason for thinking the explanation is correct, require me to know, or be in a position to know, that it is not correct? For virtually anything that could be true, there is some possible explanation of why. Is this mere possibility sufficient to undermine the force of positive grounds for belief?

So far, we have not seen adequate reason to reject the commonsense view

that we have much knowledge and numerous justified beliefs. But even if skeptical arguments do not require rejecting this view, they may prevent our showing that it is true or even providing minimal justification for it. Let us, then, explore the possibility of doing this.

## **Deducibility, evidential transmission, and induction**

When we come to the problem of induction, it seems clear that one assumption the skeptic is making is that if we believe something on the basis of one or more premises, then we can know it on the basis of those premises only if it follows from them, in the sense that they entail it. Call this the *entailment principle*. It says in effect that *knowledge can be transmitted only deductively*.

Why should we accept this principle? Not simply because inductive reasoning is “invalid”; for that term may be held to be improperly applied to it: inductive reasoning is strong or weak, probable or otherwise, but it does not even “aim” at (deductive) validity. Even if it may be properly said to be (deductively) invalid, however, that may be considered an uncontroversial technical point about its logical classification. It is a point of logic, not of epistemology. So conceived, the point does not imply either that knowledge of the premises of inductive reasoning cannot ground knowledge of its conclusions, or that justified beliefs of those premises cannot ground justified beliefs of their conclusions.

One might, on the other hand, accept the entailment principle and argue that when properly spelled out inductive reasoning can be replaced by valid deductive reasoning. For instance, suppose we add, as an overarching premise in inductive reasoning, the uniformity of nature principle, which says that nature is a domain of regular patterns that do not change over time. From this together with the premise that the sun has always risen each day it apparently does follow that it will rise tomorrow.<sup>1</sup>

But what entitles us to the premise that nature is uniform? Hume would reply that it is not knowable a priori, and that to say that we know it through experience—a way of knowing it that would depend on inductive reasoning—would beg the question against him. For on the Humean view, if our belief of the uniformity principle is grounded wholly in premises that only inductively support it, we do not know it. I believe that this Humean response is highly plausible. The problem of induction must be approached differently.

### *Epistemic and logical possibility*

What perhaps above all makes the entailment principle plausible is the thought that if our premises could be true and yet our conclusion might be false then we cannot know (or even justifiedly believe) the conclusion on the basis of those premises. At first, this thought may sound like just another formulation of the entailment principle. It is not; it is different and considerably

more plausible. That is partly why, when it is conflated with the entailment principle, it seems to support that principle. The 'might' in question is epistemic; it is like a farmer's in 'That wood dust might mean carpenter ants' or a physician's in 'Those abdominal pains might mean appendicitis'. This 'might' suggests not only that for all we know (or may take ourselves to know) the pains *do* mean appendicitis, but also that there is reason for at least some degree of suspicion that there is appendicitis and perhaps some need to rule it out.

The statement that certain abdominal pains might mean appendicitis is not merely an expression of a bare logical possibility of appendicitis—a statement that appendicitis is possible without contradiction—based, say, on no one's being absolutely invulnerable to it. If that very weak and general statement represents all we know about the case, we are not entitled to say that the pains might mean appendicitis. It is also not a logical impossibility that the Tower of London levitate above the City; but we would be quite unjustified in saying that it might.

Three notions must be distinguished here. The epistemic 'might' just illustrated suggests some reason to believe the proposition in question and might be called epistemic probability—not because the probability need be high but because it bears on knowledge and implies that, relative to all we know or are justified in believing, knowing any contrary proposition requires something like ruling out the truth of that proposition. As one would expect, epistemic probability is stronger than epistemic possibility—roughly, consistency with what we know or may at least take ourselves to know. Both are quite different from mere logical possibility—what can be the case without contradiction or some other kind of necessary falsehood.<sup>2</sup>

Epistemic possibility does not entail epistemic probability, and neither is entailed by mere logical possibility. This point bears importantly on the problem of induction. It is true that if, no matter how good inductive reasoning is, its premises could be true and yet its conclusion might, in the epistemic sense, be false, perhaps we cannot know the conclusion on the basis of them. But is this generally the case with inductive reasoning? I cannot see that it is.

Moreover, suppose it could be true that, relative to its premises, the conclusion of inductive reasoning might, in the epistemic sense, be false, what reason is there to think that this really is true? Skeptics cannot justifiably argue for this claim as they sometimes do, maintaining, simply on the ground that the premises do not entail the conclusion, that the conclusion might be false. Arguing in this way is rather like saying, of just any stomach ache a child gets after eating too much Halloween candy, that it might mean appendicitis.

It is barely possible that, relative to all we know or are justified in believing about the child, the stomach ache means appendicitis. But from that bare possibility we may not automatically conclude that appendicitis is epistemically probable—roughly, that relative to all we know or are justified in believing, we are unjustified in disbelieving that the stomach ache might mean

appendicitis. Nor does this bare possibility rule out our knowing, on inductive grounds, that overeating is the cause.<sup>3</sup>

### *Entailment, certainty, and fallibility*

There are other reasons for the attractiveness of the entailment principle, at least from a skeptical point of view. If one embraces the infallibility principle, one is in fact committed to the entailment principle. For suppose that, from known—and hence on this view infallibly believed—premises, one inductively derives a belief which is not itself infallible, as (empirical) beliefs which are inferentially grounded typically are not. Because inductive transmission allows inference of a false conclusion from true premises, the belief I inferentially derive could, as far as sheer logic goes, be false despite the truth of its inductive premises and my infallibly believing them. True premises, even if infallibly believed, simply do not absolutely guarantee the truth of a conclusion inductively inferred from them. Hence, beliefs of such inductively inferred conclusions would be fallible (unless they happened to be self-grounding or to have necessary truths as objects). But then, being fallible, these beliefs would be capable of falsehood and hence would not constitute knowledge. Thus, knowledge can be inferentially transmitted only by deductive inference. *Only valid deduction inferentially preserves infallibility.*<sup>4</sup>

If one thinks of knowledge as entailing absolute certainty, one might again be drawn to the entailment principle. For even if a fallible belief can be absolutely certain, a belief that is only inductively based on it will presumably be at least a bit less certain and thus not absolutely certain. For the proposition believed—the conclusion belief—is supported by the original belief with only a probability of less than 1 rather than with absolute certainty, as when the conclusion is entailed by the premises. This would allow that the premise belief be certain and the conclusion belief not certain (or less so), as it would not inherit from the premise belief the same degree of protection against falsehood.

To see this, suppose that the premise belief only minimally meets the standard for absolute certainty. Then a belief inductively grounded on it can fall below that standard and thereby fail to be knowledge. Putting the point in terms of probability, we might imagine a case in which our premise meets the bare minimum conditions for absolute certainty, which we might represent by a probability of 1. Then, any conclusion that follows only inductively from this premise will thereby inherit from it only some lower probability and hence fall below the minimum level for absolute certainty. Thus, again the skeptic will argue that only deduction is sufficient to transmit knowledge.

But we have already seen reason to doubt both the infallibility principle and the view that a belief constitutes knowledge only if its status is absolutely certainty. Indeed, I do not see that skeptics give us good reason to believe either these principles or the entailment principle. It does not follow from the



absence of good arguments for the principles that they are, as they appear to be, false; but if there is no good reason to believe them, even skeptics would approve of our refusing to accept them.

Absolute certainty is a high, and in some ways beautiful, ideal; but it is neither adequate to the concept of knowledge nor appropriate to the human condition.

## **The authority of knowledge and the cogency of its grounds**

There is one further principle we should consider, one rather different from those examined so far and apparently more modest. It derives in part from the idea that if you know something, you have a certain authority regarding it, an authority presumably due to your being in a position, by virtue of some ground of your knowledge, to see the truth you know. This authority is in part what accounts for the possibility of knowledge through testimony: if you know something, you have an authority about it such that normally I can come to know it, and commonly also to acquire justification for believing it, from your testimony.

### *Epistemic authority and cogent grounds*

Indeed, if you tell someone that you know something—especially when you are asked whether you really know it—you put yourself on the line. It is as if you gave your firmest assurance—an epistemic promise, as it were—that it is true. If it turns out to be false, your position is somewhat like that of a person who has broken a promise. You are open to a kind of criticism and may have to make amends. A good theory of knowledge should account for this epistemic authority that seems to go with knowing—or to be, in many contexts, implicit in attributing knowledge to someone.

A stronger but closely associated view is that if you know that a proposition is true then you must be able to say something on behalf of it. After all, the question how one knows is always intelligible, at least for beliefs that are not of luminously self-evident truths or of self-ascriptive propositions about one's current consciousness (two kinds of belief not in question for the most important kinds of skepticism); and if one really does know, one should be able to give more than a dogmatic answer, such as 'I can just see that it is true'.

The associated principle might be expressed in what I shall call the *cogency principle*: with the possible exception of beliefs of certain self-evident propositions and certain propositions about one's current consciousness, *one knows that something is so only if one has grounds for it on the basis of which one can (in principle) argue cogently for it*.

Since the cogency principle requires only that one can argue cogently for what one knows, temporary inability to mount an argument would not preclude one's knowing. Even little children might have knowledge, for perhaps

if they could just find a way to express their grounds they could provide cogent arguments. And since self-evident propositions and propositions about one's current consciousness are knowable even according to most skeptics, and may be objects of directly justified belief, there is a stopping place in epistemic chains and no regress need result when one produces a series of arguments to support a claim. What is known must simply be either traceable to those secure foundations or otherwise defensible by appeal to adequate grounds.

If the cogency principle is combined with the entailment principle, it will immediately preclude anyone's having knowledge on inductive grounds; for the entailment principle implies that inductive grounds are never cogent. But it need not be combined with the entailment principle. If it is not, it can allow for inductive reasoning of certain kinds to be cogent and thereby to transmit knowledge.

Even a moderate skeptic, however, is likely to accept at most a restricted kind of induction, a kind whose premises make its conclusion at least close to certain. This kind meets a higher standard than is usually applied to inductive inference. Thus, even though the cogency principle is separable from the entailment principle, it need not be combined with the entailment principle to be very hostile to the commonsense view that we can know the sorts of things I have been suggesting we can know, at least if this view is understood in a foundationalist framework. For this principle strikes at some of the main sources of knowledge as they are plausibly understood, and it threatens to undermine our claim to knowledge of the past, the future, and the external world. Let us pursue this.

It is true that some of our beliefs that constitute direct knowledge (and are directly justified) can be supported by apparently more secure premises. For instance, my belief that I see a green field can be supported by premises about how things appear to me, which concern only my present consciousness. After all, that this is so seems to be the only good explanation of why my visual field contains a green field. But this supportability by premises need not hold for everything it seems reasonable to regard as directly known. It may not hold for apparently memorial knowledge. As we saw in Chapter 3, one might know something through the success of one's sheer retentive powers even when the only premises one knows or is justified in using to support it fail to justify it.

A proponent of the cogency principle would certainly tend to deny that my memory can be trusted as a source of direct knowledge or direct justification, in part because memory seems far more liable to error than perception. Moreover, I might be unable to provide good inductive reasoning to support the reliability of memory even in cases in which it is very vivid, if only because such reasoning would require my depending on my memory for my justification in believing its premises, say premises about how often my past memory beliefs have been confirmed. To summarize their track record, I must remember how things turned out in the past—or at least remember

that I wrote the results down as they occurred. I would thus be relying on memory in order to vindicate it. Still, even if I could give no cogent argument to justify my memorial beliefs, it does not follow that they are not justified, or that they do not constitute knowledge.

### *Grounds of knowledge as conferring epistemic authority*

Must we accept even the apparently modest cogency principle, which requires that in order to know something, we need grounds for it from which we can (in principle) argue cogently for it? I do not see why. Certainly one can have a kind of authority without being able to defend it by premises or exhibit it in argumentation. Consider someone who can always tell “identical” twins apart but cannot say how. Moreover, ‘I see it’ need not be a dogmatic answer to ‘How do you know?’ It may simply specify one’s grounds, as when one says, ‘I see it’ in answer to ‘How do you know there is still ice on the road?’ It says *how* one knows; it need not (though it may) show that one does, particularly if showing this requires more than exhibiting an appropriate source of the challenged belief.

There is a general lesson here. When skeptics ask how we know something, this is typically a challenge to show it. I have already argued that knowing something does not require being able to show it, so this challenge is not always appropriate. What I now want to stress is that the skeptical emphasis on ‘How do you know?’ as a request for a demonstration must not be allowed to obscure the possibility of taking it as just a request to specify a ground of one’s knowledge and of fulfilling that request simply by giving an adequate ground. In doing this successfully, one shows how one knows.

In saying how I know something by citing my ground I may also be doing something further: exhibiting my knowledge of what the ground supports and perhaps even of the fact that it does so. But even exhibiting knowledge need not be showing that one has it. I can exhibit knowledge that there is black ice before me by taking my foot off the accelerator and coasting over it while saying “Tighten your seat belt”; this is not doing anything with the epistemic conclusion that I know there is ice before me. Moreover, even if I know the abstract proposition that, say, seeing ice on the roads shows that they are icy, I may not know how to use this connection to show that I know they are icy. Still, if my knowledge that they are carries an authority that can be confirmed by my citing my ground, why need I be able to go on to the sophisticated task of showing that I have knowledge? Again the analogy to virtue is pertinent: having virtue requires being able to manifest it in appropriate circumstances, but not being able to show—in any further way—that one has it.

### *Exhibiting knowledge versus dogmatically claiming it*

One might think this approach licenses dogmatism. Granted, saying ‘I see it’ could be dogmatic if intended to show conclusively that I know, for instance

by serving as absolutely proving that there is green grass before me. But the same words can simply indicate the basis for my knowledge. This is different from flatly claiming that I have knowledge. Indeed, saying it cites a ground for my belief which, if adequate, suggests that I am not being dogmatic in taking myself to know. Perhaps it is precisely because the skeptic's 'How do you know?' is commonly meant as a challenge to be shown conclusively that one knows, and not as a request to specify a source or a ground of the knowledge, that saying 'I see it' seems dogmatic in the context of discussing skepticism, even when the function of saying this is mainly explanatory.

If the issue raised by skepticism is whether we can show that we have knowledge, the point that an appeal to visual experience does not conclusively establish visual knowledge is an important concession. But the issue here is whether the skeptic succeeds in showing that we do not have perceptual knowledge. In that context, the point is not a concession. Once again, we can see how skepticism can gain credibility because skeptics make it sound as if their case against the existence of one or another kind of knowledge succeeds if we cannot show that there is such knowledge. In fact, we need not be able to show that there is knowledge in order to have it; and the skeptic must give us good reason not to believe that there is knowledge.

## Refutation and rebuttal

Have I, then, refuted skepticism, even in the few forms considered here? I have not tried to. Refutation would require showing that those forms of skepticism are false, which in turn would entail showing that there is knowledge (and justified belief). What I have tried to do is to rebut skepticism in certain plausible forms, to show that the arguments for those skeptical views do not establish that we do not have knowledge (and justified belief). Refutation of skepticism suffices for its rebuttal; but rebuttal does not require refutation. Now suppose I have succeeded in rebutting skepticism. Where do we stand? May we believe that we have knowledge, or may we only suspend judgment both on this and on skeptical claims that we do not?

I have already argued, by implication, that one can know something without knowing that one knows it. For instance, in arguing that much of our knowledge is not self-conscious, I indicated how I can know that there is a green field before me without even believing that I know this. I do not even form such self-conscious beliefs in most everyday situations. Moreover, toddlers, who do not understand what knowledge is—and so are not in a position to believe they know anything—can apparently know such simple things as that Mama is before them.

Even if I did have the second-order knowledge that I know the field is there, I surely would not possess (as seems impossible in any case) the infinite series of beliefs required by the view that *knowing entails knowing that one knows*—the *KK thesis*, as it is called—the series that continues with my knowing that I know that I know; knowing that I know, that I know that I know; and so forth. There is no plausibility in thinking that if I know that (for example)

the field is there, I must know that I know that I know . . . this, up to the limit of my capacity. I never have such a repetitive thought. Moreover, I do not think that I believe (or disbelieve) the proposition in question (I have not tested my memory here); and I cannot imagine a good use for it.<sup>5</sup> Given these points (among others), it would be a mistake to think, as some skeptics might like us to, that if we do not know that we have knowledge, then we do not.

For similar reasons, it seems possible that we might be justified in believing that we have knowledge even if we are properly unwilling to claim that we know we do, and perhaps even if we are properly unwilling to claim justification for believing that we do. Let us explore these possibilities.

If foundationalism is correct, then if one can know anything, one can know at least something directly. Moreover, some of the sorts of things that a plausible foundationalism says we know directly—for instance, self-evident truths and some propositions about our present consciousness—are the kinds of things which, simply on the basis of reflection on the examples involved, it is plausible to think we know. Perhaps, of course, this intuitively plausible point, even if it does not involve arguing from premises, shows that we have knowledge. In any case, I think that we are justified in believing that we have some knowledge even if we cannot show that we do; and I am aware of no good argument against the view that we have some knowledge even of the external world.

Might there be a way, however, to give a cogent, positive defense of common sense: to show that we have knowledge, even of the external world? And could we establish this second-order thesis even to the satisfaction of some skeptics? There is no satisfying a radical skeptic, one who denies that there can be any knowledge or justified belief (including justification of that very claim, which the skeptic simply asserts as a challenge). For a radical skeptic, nothing one presents as a reason for asserting something will count as justifying it.

Could anything be said, however, to show that there is knowledge to a moderate skeptic: one who holds, say, that although transmission of justification and knowledge must be deductive, we may justifiedly believe, and perhaps know, at least self-evident propositions and propositions about our present consciousness? Even if the answer is negative, perhaps we can show that there is knowledge, or at least justified belief, whether any skeptics would find our argument plausible or not.

### **Prospects for a positive defense of common sense**

How might an argument for a positive defense of common sense go? Let us consider justified belief first, as showing that certain of our beliefs are justified, unlike showing that some of them constitute knowledge, does not require showing that the beliefs in question are true.

### *A case for justified belief*

One might view the issue this way: if we are to show that there are justified beliefs, then one result of our argument will itself be producing justification, specifically justification for the second-order belief that there are justified beliefs. For to show something by argument entails at least producing justification for believing it.

If we are to provide such second-order justification, we apparently need at least two things: a general premise expressing a sufficient condition for justification, and one or more specific premises saying that a particular belief meets that condition. For instance, the general premise might be the justification principle that

- 1 An attentively held belief to the effect that one is now in an occurrent mental state, such as thinking, is (*prima facie*) justified,

where an attentively held belief is one based on careful attention to the proposition in question, and where the justification is not absolute but *prima facie*: it must simply be strong enough to make it appropriate for a rational person to hold the belief.<sup>6</sup> The particular premise might be that

- 2 I have an attentive belief that I am now in such a state, namely thinking.

If I am justified in believing these premises, I am justified in inferring deductively from them, and thereby in believing on the basis of them, what they self-evidently entail, namely, that

- 3 My belief that I am thinking is (*prima facie*) justified.

Here I would be inferentially justified in believing (3), at least if I can hold all three propositions before my mind in a way that avoids dependence on memory of my premises; and this seems possible for me. (If I needed so many premises, or such complicated premises, that I could not hold them in mind at once, then my justification for my conclusion would depend on that of my memory belief(s) of my premises.) Premises (1) and (2) self-evidently entail (3), and a moderate skeptic will very likely grant that if, from premises that I am justified in believing, I infer (without dependence on memory) a conclusion that self-evidently follows from them, I am justified in believing that too.

But how am I now justified in believing premises (1) and (2), if I am? There is some plausibility in holding that the general principle, (1), is justifiable directly (non-inferentially) by reflection (and is arguably self-evident), and so my belief of it might itself be directly justified. This is not to deny that it could be justified by prior premises, since self-evidence does not entail

unprovability or unevidenceability. The point is only that it is arguably justified by reflection not dependent on one's appealing to such premises. As for the particular premise, (2), I might be directly justified in holding it by virtue of a justification principle similar to the general one, but applying to beliefs (dispositional elements), a principle to the effect that if, on careful introspection, one believes that one attentively holds a proposition, *p*, then one is justified in believing one does (presumably directly justified, if one has introspected carefully).

Now if my belief of my general premise is justified, and if I may justifiedly hold the particular premise, then surely I may justifiedly conclude that I am justified in my belief that I am thinking. I may justifiedly conclude this even if my justification in believing my premises is not direct, as I am tentatively assuming it is. Moreover, if my beliefs of (1)–(3) are true, they may also constitute knowledge: my justification for each seems strong enough, and apart from this matter of degree there appears to be no other kind of bar to knowledge.<sup>7</sup>

Given the plausibility of the premises just used to try to show that I am justified in holding a belief about my own mental life, I am inclined to think that it can be shown that there are some justified beliefs, even some justified empirical ones. But even if the line of argument I have used is successful, one might question whether it extends to any beliefs about the external world. What would be our general principle for, say, visual perceptual beliefs?

In answer, perhaps we might begin with an instance of a justification principle stated in Chapter 1. Applied to the green field example, this would say that

- a If, on the basis of a vivid and steady visual experience in which one has the impression of something green before one, one believes that there is something green before one, then one is (*prima facie*) justified in so believing.

Surely we may say this, particularly since the justification in question is admittedly defeasible. (It could, for instance, be undermined by my justifiedly believing that I have frequently been hallucinating greens lately.) Suppose this premise may be believed with direct justification, and we may also believe (possibly with direct justification) that

- b I have a belief (that there is something green before me) grounded in the way the premise—principle (1)—requires.

Then I may, much as before, justifiedly conclude that

- c I am (*prima facie*) justified in believing that there is something green before me.

To be sure, my perceptual belief is only *prima facie* justified—roughly, justified in the absence of defeating factors. But this is still a significant conclusion, even if (as seems possible) I could not, by reflection alone, rule out all of those defeaters.

### *The regress of demonstration*

Supposing this line of argument against the skeptic is sound, have I shown anything? If showing something is producing a good argument for it from true premises that one is justified in believing, presumably I have. It is easy, however, to think that the skeptic would be correct in denying that I have shown anything. For there is a subtlety here that is easily missed. Even if I have shown my conclusion, I might not be justified in saying, to the skeptic or anyone else, that I *have* shown it, or even in believing I have. For justification for asserting or believing that second-order proposition about my beliefs would ordinarily require holding (or at least having justification for holding) *third-order* beliefs, such as the belief that my second-order belief that I believe I am thinking is justified and true (since this second-order belief has been shown by good argument). And what in my situation would give me that still higher-order justification?

The general point is that whatever one's justificational or epistemic achievement, justifiedly saying or even justifiedly believing that one has succeeded in it requires justification or knowledge at the next higher level. This higher-order justification or knowledge may or may not be forthcoming. Initially, this point may seem to doom my original attempt to show that I have a justified belief. But I do not think it does; it brings out only that one can show this without automatically showing the higher-order proposition that one has shown it. Plainly, we can achieve something even if we are in no way entitled to credit ourselves with achieving it.

What we have encountered here is a counterpart of the regress of justification, a *regress of demonstration*: if one shows anything at all and can be asked to show that, then there will be either an infinite regress of demonstrations, or a circle of them, or there will be some unshown shower (at least one unshown premise). In a sense, there will always be some point at which self-congratulation—or a final dismissal of the skeptical challenges—is inappropriate. This is another reason why a moderate foundationalist perspective should not be thought to lead to dogmatism.

Still, even if I do not know that I have shown that my belief that I am thinking is justified, I may yet have shown this; and if I have, then I may well know the proposition that I have shown: that my belief that I am thinking is justified. Perhaps, moreover, a similar procedure may be repeated, with equal success, at each higher level to which one can ascend without losing track of the progressively more complex issue. Then, with sufficient patience, one could show any given justification-ascribing proposition in the hierarchy—



that one has shown that one has a justified belief, that one has shown that one has shown this, and so on, to the limit of one's comprehension.

### *A case for knowledge*

The reasoning we have been exploring in connection with justification also bears on skepticism about knowledge. If the premise beliefs, (a) and (b), are true, they may constitute knowledge. I may, then, not only be showing that I am justified in holding a belief about the external world; I also may, as a result of my reasoning, know that I am justified in holding it. There would then be at least justificational self-knowledge: self-knowledge to the effect that one has justification for some beliefs.

In reasoning from (a) and (b) to (c), however, I do not automatically know that I am showing that I am justified in believing something about the external world. Suppose I do not know this. Perhaps I only hope that I am showing it. Then, even if I do have second-order knowledge that I have a justified belief about the external world, I may not be justified in holding the third-order belief that I have (second-order) knowledge that I have this (first-order) justified belief, the belief that there is something green before me. I have as yet no principle that would justify me in concluding that I know or justifiedly believe that I have a justified first-order belief. I lack a principle stating conditions that generate second-order knowledge or second-order justified belief.

It seems, however, that the sort of justification I apparently have for all the relevant beliefs, including the belief that I have a justified belief about the external world, is the kind whose possession by true beliefs is sufficient for their constituting knowledge. Thus, through reasoning using premises such as (a) and (b) I may well know that I have justified beliefs about the external world. Certainly I have reason to think that the skeptic does not know, or justifiedly believe, that I lack justified beliefs about the external world.

One assumption it is natural to make in using this strategy against skepticism deserves emphasis: the assumption that the crucial principles of justification are a priori, and believing them is justified by reflection (directly, or at least on the basis of self-evident steps from directly justified beliefs of a priori premises). Suppose the principles are empirical. Then our justification for believing them would presumably be broadly inductive. A skeptic could plausibly deny that, on an inductive basis, we can justifiedly believe them. Let us pursue this possibility.

### *A circularity problem*

There would apparently be a circularity problem if we had to justify our crucial principles inductively. For justifying them by inductive reasoning would seem to presuppose using just such principles, principles that specify, for instance, under what conditions inductive inference can transmit

justification or knowledge. We would have to use induction to develop a track record for such inferences, say by determining, observationally, how often their premises are found true and their conclusions determined to be false. We would need to rely on perception and memory to do this—in addition to using induction to infer from a good track record on the part of a source to its general reliability. To acquire justified beliefs about the reliability of perception, moreover, we would need to use perception, say in looking at objects again to see if our initial color perceptions were accurate.

Are the kinds of principles of justification I have been using *a priori*? That is certainly arguable; but it is also controversial. On the most plausible kind of reliability theory of justification, for instance, a belief is justified by virtue of being grounded in reliable belief-producing processes such as perceptual ones; and it is apparently not an *a priori* matter what processes are reliable, that is, actually produce a suitably large proportion of true beliefs. This is the sort of thing that must be determined largely by observation.<sup>8</sup>

Thus, for reliabilism about justification, in order to know what principles account for justification, one must know what processes tend to generate true beliefs. One could determine that only through considerable experience. Hence, if these principles are empirical, the circularity problem just mentioned would beset the attempt, within a reliabilist framework, to justify them. We would have to use perception to vindicate perception, for instance checking on our visual perceptions by relying on tactile ones. Our conclusion could of course be true, but to rely on perception in arguing that it is reliable would not justify one's conclusion. It is an interesting question how much the circularity may be mitigated insofar as, in confirming the reliability of one sense, we use (as just illustrated) only the deliverances of a different one. This widens the circle, but it does not provide a way out of it. Indeed, to validate any source over time, we must rely on memory to track our successes and failures. Without memory, our lives would be a series of disconnected snapshots.

On the other hand, I argued above that reliability theories are less plausible for justification than for knowledge, and I believe that it is more reasonable, though by no means obviously correct, to suppose that at least some principles about the conditions for justification are *a priori*. I would include various principles expressing ways in which—as described by Chapters 1 through 5—justification is produced by its basic sources.

## **The challenge of rational disagreement**

As so far described, skepticism applies to much that we ordinarily take ourselves to know or at least to believe quite justifiedly. But skepticism may be both more selective and more worrisome even when less comprehensive—when it takes a smaller bite. It may plausibly arise when we disagree with others we consider generally credible on the topic in question, say the existence of God or of objective moral truths. This example shows both how

skepticism may be highly selective and how it may be disturbing. Skepticism about such matters is highly selective because some people seldom if ever encounter such disagreements; and it is disturbing because it challenges the rational basis of beliefs that guide important segments of our conduct.

### *Intellectual pluralism*

This is an age of intellectual pluralism: there is a great diversity of views on many important matters. It is also an age of rapid and widespread communication: through various media and, especially, the internet, more and more people are aware of views opposing their own and, often, of evidences for those views. Should we respond skeptically and conclude that in the face of such widespread disagreements, we rarely have knowledge or justified beliefs, at least in controversial matters?

The mere presence of widespread disagreement on a topic, say the objectivity of ethics, does not imply that no one has knowledge of one of the disputed positions. Nor does that skeptical conclusion follow even when one party has more evidence than the other(s). *Quantity* of evidence does not guarantee a *quality* sufficient to justify the proposition the evidence supports. This point has been disturbingly neglected by political advertisements that say such things as ‘Twenty million Americans can’t be wrong’. Even on the assumption that their testimony can provide *some* evidence, this claim is simply false. All of humanity has been wrong on such matters as the shape of the earth.

Moreover, as we have seen in earlier chapters, some beliefs—such as some beliefs of self-evident propositions—may be justified on an intuitive basis that is not, in everyday parlance, viewed as a kind of evidence. The intuitive sense of truth underlying such beliefs as that nothing is round and square—an intuitive seeming, in one terminology—may be considered a supporting *ground* even if it is not considered a kind of evidence. But people have conflicting intuitions as well as conflicting beliefs based on ordinary kinds of evidence. Cognitive pluralism is not limited to belief, much less to theories or other complex views. To see how a skeptical problem is presented by rational disagreement in cognitive matters broadly construed, we must refine the issue.

### *Epistemic parity*

The first thing to note is that dealing with rational disagreement presupposes a measure of non-skeptical confidence. Indeed, if I did not trust my experience in the first place, what would justify my belief that others differ from me and, despite this, meet a standard of rationality I respect? It turns out, surely, that to assess challenges to our views we must—whether in religion, science, philosophy, or other cognitive domains—consider *both* our experience and

our overall intellectual position, including our resources for explaining why others disagree, in a way that preserves as much as possible of our view.

More specifically, it is at best difficult—especially in complex cases—to acquire justification for believing someone else to be an *epistemic peer* with respect to some proposition, *p*, on which the other person disagrees. Roughly, to be an epistemic peer in a given matter is to be (a) exposed to the same relevant evidence as oneself, (b) equally conscientious in considering it, and (c) equally rational in the matter. Even beyond this, there is the question whether the other person might have background beliefs—such as a belief that Humean skepticism about induction is sound—that, perhaps by introducing biases, reduce the person's overall justification regarding *p*.

At least three further factors should be noted. One is that someone else's disbelieving *p* is itself a reason (even if not a strong one) for a person who *rationaly* believes *p* to doubt that the other is a full-scale epistemic peer in the matter.<sup>9</sup> The second is that we are better positioned to make a critical appraisal of our own evidence—at least when it is experiential, as with memory impressions and intuitions—and of our responses to it, for instance in assessing whether our belief that *p* is based on the evidence rather than on, say, wishful thinking, than of anyone else's evidence or responses to that evidence.<sup>10</sup> Other things equal, then, we are better justified in our assessment of our own basis for believing *p* and of our response to that basis than in our assessment of the basis of anyone else's believing it or of anyone else's response to that basis.

A third factor is this. As we check and re-check our own grounds for a justified belief that *p* and our responses to them, we tend to *increase* our justification for believing *p*, at least when we retain that belief *in the light of this effort*. Indeed, even if, from a skeptical disposition, we do not retain it, our propositional justification regarding *p* may still rise, as when we come closer to having conclusive evidence but, being highly cautious regarding *p*, withhold belief from it.

### *Dogmatism, fallibilism, and intellectual courage*

These considerations provide some support for the modest conclusion that the very exercise of critically seeking to establish the epistemic parity of a disputant may give a rational person a justificatory advantage in the dispute. Perhaps we may conclude that other things being equal, a rational conscientious attempt to establish the epistemic parity of a disputant tends to favor the conscientious inquirer who, on the basis of such an attempt, retains the belief that *p*. This conclusion, qualified as it is by a tendency element, does not entail that such conscientious inquirers should never withhold *p* instead (neither believing nor disbelieving it). That might be a reasonable thing to do in some cases, especially if one justifiedly believes a peer has made a conscientious attempt of the same kind and retained a contrary belief.

Where conscientious inquirers should not withhold  $p$ , retention with humility might well be a commonly justified result of the kind of reflection in question. Such retention, moreover, should normally be accompanied by a sensitivity to, and in some cases a search for, further evidence concerning  $p$ . Dogmatism, then, as implying a belief or strong disposition to believe that one is correct and the disputant is wrong, or a disproportionately strong belief of  $p$ , is inappropriate in all cases of rational disagreement between peers.

To be sure, each of two disagreeing peers can go through the same kind of conscientious review of the apparent parity of the other, in which case the enhancement in justification may be equivalent on each side. Moreover, repeating such a review should not be expected to increase one's justification for  $p$  indefinitely and will tend to do so decreasingly with each repetition. Given these qualifications, it should be stressed that the kind of increase in justification that such critical reflection may bring does not automatically warrant reduced humility and may indeed call for more of it.

It turns out, then, that sometimes our plausibly challenged beliefs should, on balance, be retained; sometimes our *unchallenged* beliefs should be given up. There is no simple formula here, in ethics, philosophy, religion, or any other domain. A good attitude with which to approach rational disagreement with someone we consider an epistemic peer in the matter is fallibilism of the kind described in Chapter 12. This attitude contrasts with dogmatism but is consistent with intellectual courage. That courage is among the intellectual virtues. It is roughly a settled disposition to form beliefs, and to retain one's existing views, with the right level of conviction in relation to the strength of one's grounds, and with a kind of fallibilism that leads one to countenance disagreement but not to give up one's views under the mere pressure of widespread disagreement.

## Skepticism and common sense

Where, then, does this chapter leave us with respect to appraising skepticism? To begin with, there are forms of skepticism I have not mentioned, and I have also not discussed every plausible argument for the skeptical principles I have addressed: chiefly the infallibility, certainty, back-up, entailment, and cogency principles. But these principles are in some important ways representative of those on which even moderate skepticism rests. I have offered reasons for rejecting them, and on that basis I have maintained that skepticism, at least insofar as it depends on these and similar principles, can be rebutted. It can be shown to be rationally resistible.

We are, then, warranted in refusing to accept skepticism concerning justification and knowledge of propositions other than those that are self-evident or self-ascribe to the believer a present occurrent mental property. If it is not false, it is at least not justified by what seem the main arguments for it. It is not clear, however, that anything said above refutes the kinds of skepticism we have considered. For refuting those views entails showing them to be false, and it is not altogether clear what that requires.

Positively, I have suggested that on one plausible notion of showing something, namely, validly and justifiably deducing it from true premises which one justifiably believes and are good grounds for it, we can show that there are some beliefs we are justified in holding, probably even some justified beliefs about the external world (these may perhaps include some about the inner lives of others).<sup>11</sup> I am less inclined to say that we can—by this strategy—show that there is knowledge, particularly knowledge of the external world. Much depends on the kind of grounding required for such premises as that I see a green field before me, which, because simple seeing entails the existence of the object seen, in turn entails that there is something external. Much also depends on how rigorous a standard of showing is appropriate.

On balance, then, I have supported the commonsense view that we can know that there is both justified belief and knowledge about the external world, and can know this even if we cannot show that there is. I also maintain that there is justified belief and knowledge about one's own consciousness and about certain *a priori* matters. Skeptics certainly do not seem to have shown that we do not have knowledge and justified belief of these kinds. I believe that we have both.

Moreover, if, as argued in Chapter 12, it is true that rationality is a more permissive notion than justification, then whatever the anti-skeptical case for our having justification, it will count more strongly for the counterpart views concerning the rationality of our beliefs and other epistemic attitudes. Even if rationality, as applied to beliefs, is significantly weaker than justification, it is still the kind of status skeptics tend to deny is ever achieved by our beliefs about the external world, the past, and many other things.

Perhaps viewing knowledge, justification, and rationality in the way I have might be thought to be an article of epistemological faith. I do not think it is; but the difficulty of determining whether it is partly an article of unverifiable faith, or can be established by cogent argument, or is more than the former yet less than the latter, is some testimony to the depth and complexity of skeptical problems.

## Notes

- 1 It does not strictly follow unless we define regularity to preclude the following kind of thing: the sun rises every day except every trillionth after the Earth came to be, where tomorrow *is* the trillionth.
- 2 As noted in Chapter 5, the synthetic *a priori* might be necessary without being logically necessary in the strict sense; similarly, something can be a synthetically necessary falsehood (thus impossible, as is a round square), but not strictly logically impossible.
- 3 Epistemic possibility for us is sometimes characterized simply in terms of what is possible given what we *know*, but I think it is appropriate to include justification here, if only because we may be better able to tell by reflection on our current overall cognitive resources what we are justified in believing than what we know.

- 4 This should be taken to apply to non-formal validity: infallibility would be preserved by the inference from something's being round to its not being square, but this is a case of synthetic a priori, as opposed to formal, entailment.
- 5 Reasons for doubting that we should posit such beliefs are given in my 'Dispositional Beliefs and Dispositions to Believe', *Noûs* 28, 4 (1994) 419–34.
- 6 We might perhaps think of the degree of justification in question here as knowledge-sufficing: the kind sufficient to render a true belief knowledge, in the absence of the sorts of cases discussed in Chapter 10 showing that justified true belief need not constitute knowledge.
- 7 I ignore here the point that I might have only situational justification for my conclusion if my believing it is not *based* on my believing my premises. Note 11 comments on this problem.
- 8 For extensive discussion of the kind of circularity problem in question here see William P. Alston, 'Epistemic Circularity', *Philosophy and Phenomenological Research* 47 (1986), 1–30. Cf. Ernest Sosa, 'Philosophical Scepticism and Epistemic Circularity', *Proceedings of the Aristotelian Society Supplementary Volume* 68 (1994), 263–90, followed by a response from Barry Stroud.
- 9 My thought here is that if you have some reason or ground for believing *p*, and you adequately understand the obvious entailment of *q* by *p* (as when *q* simply is not-*p*), then you have some ground or reason, even if not as good ground or reason, for believing *q*. To be sure, to have some ground or reason for the higher-order belief that you have this ground or reason you may well need justification for the higher-order belief that you do in fact have ground or reason to believe *p*. But this justification is commonly achieved in cases of or approaching peer disagreement. Detailed discussion of the epistemology of peer disagreement, including references to recent literature, is provided in my 'The Ethics of Belief: Rational Disagreement, Intellectual Responsibility, and Ethical Conduct', in Quentin Smith (ed.), *Epistemology: New Essays* (Oxford: Oxford University Press, 2008).
- 10 The reference to evidence here must be taken to designate grounds of an internal kind, such as the "evidence of the senses." For evidence conceived as publicly accessible supporting fact, I am not suggesting that any one person is necessarily in a better position than another to appraise it, though we may still have a kind of intrinsic advantage in appraising our response to it. But for assessing rationality the central concern is the person's experience, memory impressions, reflections, and other internal elements.
- 11 I have not directly argued that there are *justified beliefs*. For I have not argued for the premise, apparently needed for his conclusion, that we are non-inferentially justified in believing that the relevant beliefs, such as the belief that there is something green before me, are *based* on the

visual impressions constituting one's grounds for it. This basing is partly causal, and skeptics are likely to argue that justification for attributing causal propositions requires inductive, hence inferential, grounds. This is not self-evident, and I have challenged it in 'Causalist Internalism', *American Philosophical Quarterly* 26 (1989), 309–20, reprinted in my *The Structure of Justification* (Cambridge: Cambridge University Press, 1993). In any case, even if I am not justified in believing that my external world belief is based on a sense impression, but only that I *have* the relevant impression, I am justified in believing that I *have* this justification for the external world belief: we might say that I am entitled to hold it even though I may not hold it on the basis of my entitlement.





## 15 Conclusion

Once again, I look at the grassy green field before me. A haze has obscured it, but the afternoon sun now streaming down on it puts it in clear view. Its shape and its shades of color are plainly in my sight. The birds are still singing. As I look at the spruce tree, I remember cutting the poison ivy vine from it. I recall its furry stem, and the recollection is so vivid that as the scene fills my consciousness, I seem almost to re-enter the past.

I cannot help having these experiences of color and shape and sound unless I deaden my senses to the world; and normally I cannot help forming beliefs as I perceive the world or look into my consciousness. I can walk away and change the external sources of my belief. But as we saw in exploring perception, we cannot entirely resist those sources. If my senses are open to my surroundings, I perceive them; if I perceive them, I tend to form beliefs about them, even if such perception does not entail forming them: beliefs about their colors and shapes, their sounds and scents and textures, and, if they have found a place in my memory, sometimes about their history.

These beliefs seem to arise directly from perceptions and through the recollections that surface from memory; they do not emerge by a process of inference from anything else I believe. I realize that such beliefs are fallible, and I understand the profound inclination toward skepticism that we can experience as we reflect on the significance of that fallibility. Still, I find no reason to doubt these everyday beliefs, and I am convinced that for the most part they are justified and constitute knowledge.

Our beliefs are countless and varied. A vast proportion of them are stored in memory, though beliefs do not originate there. As Chapter 2 brought out, memory preserves belief, but does not by itself normally produce it. It also preserves, but does not create, knowledge. Once I come to know through perception that the spruce is taller than the maple, I may know this from memory even when I have forgotten my evidence for it.

By contrast with its preservative role in relation to knowledge, memory is a basic source of justification. It can be the only present source of justification for many beliefs stored therein, beliefs whose original grounds are long forgotten. And, like sensory experience, memory can be deceptively generous in supplying justification. Just as sensory experience can mimic

perception and thereby justify false beliefs about the external world, memorial experience can mimic genuine recollection and thereby justify certain false beliefs which have the appearance of being memorially retained. I may later discover that despite their apparent memorial authenticity, they arose from wishful thinking.

These and other cases of justified false memory beliefs show that our justification for believing what is, so far as we can tell, grounded in memory is defeasible. Nonetheless, it is significantly strong. If we had good reason to believe it lacked a certain minimal strength, we would not be justified in trusting our memories without external evidence of their reliability; and it is doubtful that we could get enough such evidence if we could not trust our memory directly in at least some cases. To test my memory of the texture of the maple by going closer to it, for instance, I must retain the belief whose truth I am trying to confirm. If memory were not a basic source of justification, we could never have a large enough store of justified beliefs to yield premises adequate for significant deductive and inductive extension of our justification. The scope of our justified belief would then be drastically narrowed; and at least a great deal, and perhaps all, of our knowledge of the past, the future, and general empirical propositions would also be undermined.

If Chapter 1 concerns what might be called outer perception, Chapter 3 explores what is sometimes called *inner* perception. When we look into our own consciousness, we find beliefs also arising in the same natural, seemingly irresistible way in which they arise from outer perception. We have, however, far more control over the scenes and events that we experience only inwardly. I can blot out my sensations of color and shape only by closing my eyes, but I can dismiss my memorial image of the poison ivy vine at will and, just as directly, call up an image of the friend who helped me pull it down. In this respect—in relation to the will—beliefs, even about elements in present our own consciousness, are more like perceptions than like images of memory and imagination. Even though I can at will cease imaging the vine, once I focus on this memory image, I cannot help believing that it represents something with a greenish cast, nor could I have come to believe this at will. The inner world, like the outer world, produces certain beliefs directly and irresistibly. And these beliefs tend to be justified and, often, to constitute knowledge.

If our only sources of knowledge and justification were perception, whether inner or outer, and memory, we would be at best impoverished. We can also turn our attention to abstract matters, even while our senses bombard us with impressions. Looking at the spruce and the maple and then further to the right where there is an apple tree at the side of the field, I realize that they are so far apart that I cannot see by direct comparison whether the spruce is taller than the apple tree. But I can see that it is taller than the maple and that the maple is taller than the apple tree. Clearly, then, the spruce is taller than the apple tree. As stressed in Chapters 5 and 6, reason makes it obvious—and it is indeed intuitively obvious—that if the spruce is taller than the maple and the maple is taller than the apple tree then the spruce is taller

than the apple. This belief is as natural, and would be no less difficult to resist when I vividly consider its propositional content than my belief that there is something blue before me when I squarely see the spruce. Clearly, this *a priori* belief is also justified, and it constitutes knowledge.

The experiences and reflections of individuals, even when well-preserved in their memories, are limited. There is a great deal we cannot know about the world without relying on others. Much of what we justifiably believe and much of what we know is, in a sense, socially grounded: based on what others have said to us, whether in person or impersonally in their writings. Testimony, the central concern of Chapter 7, is a special source of justification and knowledge. It often yields direct belief; yet unlike perception, it is not a basic source of belief. It does not yield belief apart from perception, since it must be received perceptually, as in listening or reading.

Testimony is, however, like memory and unlike perception in being unlimited in the scope and subject matter of the propositions we can learn from it. To be sure, we might gain no justification from testimony if we had none deriving from a basic source. But we would have, at best, far less justification than we do if our only knowledge and justification came from basic sources, or even from basic sources and inferences from the propositions we believe on the basis of those sources.

Once we have beliefs directly grounded in one of the five common sources of non-inferential knowledge and justification—perception, memory, consciousness, reason, and, secondarily but indispensably, testimony—we are in a position to extend whatever justification and knowledge we then have. Take a simple case. On the basis of my beliefs that the spruce is taller than the maple and that the maple is taller than the apple, together with my belief that if those things are so then the spruce is taller than the apple, I infer that the spruce is taller than the apple. I began with non-inferential beliefs grounded directly in basic sources of knowledge: perception and reason. By a spontaneous deductive inference, I extended both my knowledge and my justification. And when, on another occasion, I heard rapid knocking, believed it to sound like that of a woodpecker, and inferred that there was a woodpecker nearby, I extended my knowledge by inductive inference. Chapter 8 indicates how knowledge and justification can grow indefinitely in these ways. Inference has a virtually unlimited capacity to extend our outlook.

A picture has emerged. We are in almost constant interaction with the world, external and internal. We are regularly bombarded by sensation, often immersed in the stream of our consciousness, and sometimes occupied with the testimony of a friend or with our reflection on abstract matters, such as questions of philosophy, mathematics, or science. Beliefs are a natural product of these engagements. They arise in perception, introspection, reflection, and testimony; they are preserved in memory; they are multiplied by inference. Many beliefs are grounded in the basic sources, or preserved, as non-inferential beliefs, in memory; many other beliefs are inferentially grounded in these direct ones, in the ways detailed in Chapter 9.

This picture portrays two interconnected structures. One is constituted

by foundational beliefs anchored in the bedrock of experience and reason, whether directly or through testimony that ultimately rests on it as well. The other is a superstructure of vast complexity built from the foundations by the building blocks of inference. The theory associated with this picture of our beliefs in relation to the world is psychological foundationalism. The picture is natural; and there is much to be said for the theory.

Once our psychology is understood in this structural fashion, another natural picture, similar to the first, emerges. The theory associated with this structural picture is epistemological foundationalism. I know that the spruce is taller than the apple tree. I know this on the basis of knowing that the spruce is taller than the maple and the maple is taller than the apple tree, together with the proposition that if this is so then the spruce is taller than the apple tree. I know that conditional proposition directly, through rational comprehension of it, and I know the other premise by sight. These two items of knowledge are foundations of my (non-basic) knowledge that the spruce is taller than the apple tree.

I do not readily see how to go any further in grounding my knowledge here; and even if I can go on, it is not clear how I could have knowledge at all if there were not some point or other at which my belief is connected with the reality in virtue of which it is true: the trees with their woody skeletons and colorful foliage; the apparently unchanging abstract relations grasped by reason.

Metaphorically, this epistemological picture portrays both knowledge and justification as grounded in looking and thereby in seeing. Perception looks outward, and through it we see the physical world. Memory looks backward, and through it we see the past, or at least some of our own past. Introspection looks inward, and through it we see the stream of our own consciousness. Reason looks beyond experience of the world of space and time, and through it we see concepts and their relations. Testimony, our chief social source of knowledge, looks to others and thereby draws, however indirectly, on all of these individual sources in those who convey their knowledge to us. Testimony enables us to see—though at one remove, through the attester's eyes—virtually anything that an accurate person attests to. By attending to testimony we can look through any of the basic sources of knowledge and justification, as they have informed others, upon any subject matter they accurately describe to us.

The foundational pictures, both in epistemology and in psychology, have their appeal; yet we can imagine going further in the process of justification than they suggest we should. It may be natural to think that, at any given time, a chain of justification or knowledge will be anchored in the bedrock of experience or reason, just as its constituent beliefs apparently are—though, to be sure, the chains may be interconnected and, where one of them puts stress on another, broken or torn from their moorings. Coherentism challenges the foundationalist picture. Its proponents may grant that the picture

fits our psychological make-up. But they view the structure of our knowledge and justified belief differently.

Coherentists see the structure of our knowledge and justified belief as something like a vast fabric of interlocking fibers. Some of these may be connected to experience, but those are not privileged in generating knowledge or justification. True beliefs constitute knowledge when they are suitably woven into the whole fabric, which, in turn, must hold together in a systematic way. Justification is also a matter of how beliefs are connected with the rest of the fabric. A belief that is a largely isolated strand, for instance one that is not inferentially based on any other belief or even significantly connected with any other in subject matter, would not be justified.

This and other coherentist pictures can also have powerful appeal, particularly in understanding the process of justification, in which we commonly try to show that a belief is justified by connecting it with others that support it and thereby cohere with it. But the process of justification should not dominate our understanding of what it is for a belief simply to be justified—to have the property of justifiedness. Moreover, when it comes to knowledge, which entails truth, the coherence picture is less plausible. For indefinitely many fabrics can have internally coherent patterns; and coherentism—unless alloyed with foundationalist elements—does not require that any of the strands be anchored to the world, whether in perception or introspection or in any other way. Why, then, should we expect a coherent set of beliefs to contain truths that represent the world, or when it does, to embody knowledge?

Indeed, if a belief's being justified counts in some way toward its truth, then why should coherence alone be the basis of justification, given that coherence by itself implies nothing about truth? Furthermore, self-evident propositions, say that if no vixens are males then no males are vixens, seem such that we need only understand them to be able to know or justifiably believe them. Even if possessing this knowledge and justification implies *having* a measure of coherence in one's cognitive system, how might knowledge or justification, whether of this kind or any other, *derive* from coherence? There are plausible attempts to provide answers, but I have seen no clear success in doing so. Whether the structure of knowledge is foundational or not, I may know such things as that there is a cold glass in my hand and that there is rapid knocking nearby.

But what is knowledge?

This dauntingly simple question is the focus of Chapters 10 and 11. My knowing that, say, there is rapid knocking may seem to be simply my justifiably and truly believing this. But it is not. Through some remarkable coincidence, I could be hallucinating such a knocking while my ears, quite unbeknownst to me, are temporarily blocked. Given such veridical hallucination, I could have a justified true belief which is not knowledge. The suggested account of knowledge as justified true belief is, then, too broad. It

also seems too narrow. For there might be knowledge without justification, as with someone who, by virtue of a stable cognitive capacity, unerringly computes difficult arithmetic results with lightning speed, but is unaware of the success and is not (at first) justified in believing the answers.

We can strengthen our requirements on justification to deal with the true belief based on hallucination, and we can weaken them to deal with the arithmetic success of the lightning calculator. But it is not evident that this strategy will yield a correct and illuminating account of knowledge. We can bypass the concept of justification as a central element in understanding knowledge and try to account for knowledge by appeal to the notion of reliably grounded true belief. But it is not clear that this approach will fully succeed either, and it certainly leaves us with the problem of explaining why justification, which need not figure in the approach, has the close connection to knowledge which it apparently does have.

Moreover, justification is epistemologically important in its own right, and reliability theories seem less likely to succeed in accounting for justification than for knowledge. This is at least in part because the grounds of justification seem internal in a way the grounds of knowledge, or at least some of them, do not. We may, however, say at least this: that knowledge is true belief based in the right way on the right kind of ground. Justification or reliability or both may be essential to adequately filling out this idea; and although it is not clear just how it is to be filled out, many of the important elements can be gathered from what we have seen in this book concerning the sources, development, structure, and analysis of knowledge.

However we analyze what knowledge is, there remains the question of how much of it, if any, we have. The question is particularly important in its bearing on the major domains that concern Chapter 12: the scientific, the ethical, and the religious. It is sometimes thought that we have a wealth of scientific knowledge, as well as knowledge of certain moral principles and some knowledge of religious truths. But if what passes for scientific knowledge is often not, strictly speaking, true—or might be utterly rejected in the future—may we really say that there is scientific knowledge? If moral principles should turn out to be neither clearly grounded in experience nor plausibly regarded as *a priori*, on what basis might they be known? And if, as many philosophers think, there are no cogent arguments for God's existence and, in addition, God is not directly knowable through the experiential or rational sources that ground knowledge, how can there be knowledge of God?

These questions are very difficult. But we are warranted in giving some partial answers to them. Consider the scientific, moral, and religious domains in turn.

First, although some of what is termed scientific knowledge is no doubt mistakenly so called because it is far from the truth, there may be some propositions that are scientifically known though not precisely true, and in any case we may speak of approximate scientific knowledge in which

the proposition in view is not precisely true, but also not grossly inaccurate. Moreover, perhaps we may sometimes speak unqualifiedly of scientific knowledge, even if this knowledge is only of approximations. We may apparently speak so when the truth known is not precisely formulated, but holds within the limits of its intended application. The degree of inaccuracy within which we may speak in these ways is not sharply specifiable. But particularly when a scientific proposition yields true predictions, helps to explain other apparently true propositions, and approximates a more accurate, true proposition, we may be justified in thinking we know it and correct in calling it approximate knowledge. However we describe scientific knowledge, it should not in general be considered to result from proof, and indeed 'proof' is misleading if applied to propositions for which our evidence is inductive, as scientific evidence typically is. This point calls for fallibilism as an attitude but not for skepticism as a stance.

Second, although even our most plausible moral principles are neither obviously knowable a priori nor obviously knowable empirically, rationalist and empiricist conceptions of moral knowledge are each defensible for some moral principles. Indeed, neither conception of moral knowledge has been refuted despite sustained ingenious attempts to discredit them. Moreover, once we cast aside the common stereotype of scientific knowledge as representing chiefly a body of facts and laws discoverable by simple inductive generalization, or rigorously provable by observations, the contrast between well-confirmed scientific beliefs and reflectively grounded moral beliefs appears less sharp. It now becomes far more difficult to discredit the view that there is moral knowledge by unfavorably contrasting moral beliefs with scientific ones, as if scientific generalizations were straightforwardly "hard" and moral principles unalterably "soft." We should not conclude, then, that there is no moral knowledge. Indeed, a plausible case can be made for the view that there are some basic moral principles that are intuitively knowable and broadly self-evident. Even apart from that, it is far from clear that we should adopt either moral skepticism about moral statements or the non-cognitivist view that there are no moral propositions to be known or justifiedly believed in the first place.

Third, in the religious domain the possibility of knowledge and justification may also be defended. Even if it is true that no argument for theism is decisive, it should be remembered that a diverse group of independent but individually inconclusive arguments may, if they are mutually supporting in a certain way and if each provides some degree of justification for a conclusion, together justify that conclusion even if none by itself does. It has not been established that this point could not apply in the case of arguments for the existence of God. In any event, discussions of the question of justified religious belief and possible religious knowledge should not simply assume the evidentialist view that such propositions can be known or justifiedly believed only inferentially, on the basis of further beliefs expressing evidence for the theistic propositions in question. One or another kind of religious experience



might provide non-inferential grounds of justification, or of knowledge, or of both, somewhat in the way familiar kinds of perception do.

There are, to be sure, many important differences between the religious and perceptual cases, and what we have seen does not show either that there is or that there is not some direct knowledge of theistic propositions or direct justification for believing some of them. But apparently, even if there cannot be directly justified beliefs of them, there could be direct knowledge of them. Even supposing, however, that there could be neither theistic knowledge nor justified theistic beliefs, there might be *rational* theistic beliefs: beliefs a rational person with a certain range of experience may hold, even without having the kinds of specific grounds, or as strong grounds, as are needed for justification. The case for the possibility of rational theistic beliefs, then, is stronger than the case for the possibility of their justification. Moreover, not all religious faith entails belief. Faith that *p* (say that God is sovereign in the universe) may be non-doxastic: it does not entail belief that *p*, though it is inconsistent with disbelieving it. When non-doxastic faith is in question, it is especially important to see that the baseline for rationality is both more permissive than for justification and more permissive than for rational belief with the same content as the faith. On two counts, then, skepticism in religious epistemology needs a stronger case than is usually thought: one strong enough to defeat grounds for rational, as distinct from justified, cognition—whether faith or belief or some other propositional attitude—and far-reaching enough to undermine non-doxastic faith as distinct from belief-entailing faith, which requires stronger grounds.

Powerful skeptical arguments threaten the view that our knowledge includes as much as Chapter 12 suggests. They threaten even the common-sense view that we have some knowledge of the external world, the past, the future, and the inner lives of others. Chapters 13 and 14 are devoted to appraising some major skeptical arguments. When we realize that our beliefs concerning these domains are clearly fallible, we can begin to appreciate skeptical views. Even our sense that, whether or not we have knowledge, we do have justified beliefs weakens if we take seriously the possibility that what we accept as justification is no final guarantee of truth. But the common skeptical commitment to the ideal of infallible belief as central to knowledge is not warranted by careful inquiry into the nature of knowledge.

Infallibility may be a reasonable ideal for proof, conceived as decisively demonstrating a conclusion from rock-solid premises, such as self-evident truths or, on the empirical side, propositions about the believer's immediate consciousness. For one cannot decisively prove something—or demonstrate it, to use a term that sometimes has wider scope—from insecure premises, or by making merely inductive and hence fallible steps from even the most trustworthy premises. But why should proof be our standard of the kind of justification (or even for a kind of certainty) appropriate to knowledge? We are not talking about what is required to show conclusively that there is knowledge, but about whether there in fact is any.

If, however, we think that there is knowledge, and the skeptic challenges us, we want to show that there is. But we must not confuse—or allow skeptics to confuse—the requirements for showing that there is knowledge with the requirements for the existence of it. Perhaps it can be shown that there is knowledge. Certainly, if we want to argue for this, we need not accept the idea that showing something to be true is equivalent to proving it. But even if it cannot be shown that there is knowledge or justified belief, it does not follow that there is none. It also does not follow that we do not have something less difficult to achieve than justified belief, though significant in most of the same ways: belief that is rational.

But there surely is knowledge and justified belief. I justifiably believe, indeed I know, that that green field lies before me. Those bird songs are not fantasy. My stream of thoughts is in unmistakably clear focus. Even my recollection that I cut a vine from the spruce tree is clear and steadfast. I am justified in believing that I did, and surely I know this. We all have a huge store of beliefs of these and other kinds, including countless beliefs originally formed through testimony. These beliefs form a structure of great complexity, with innumerable changing elements that reflect our continuing experience and thought, our actions and emotions, our learning and forgetting, our inferring and accepting, our revising and rejecting, our speaking and listening. That structure is grounded in us: in our memories, our habits of thought, our mental and perceptual capacities, our rational nature.

Knowledge of the truths of reason arises within the structure itself, once we have the needed concepts. Through our consciousness of what is inside of us, and our perceptual engagement with what is outside of us, with the social world as well as our physical environment, this structure is anchored, both internally and externally, to the world. That vast and various reality is at once the ultimate source and the object of our empirical knowledge.



# Short annotated bibliography of books in epistemology

- Alston, William P., *Epistemic Justification*, Ithaca, 1989. A collection of papers centering on justification and knowledge in a way that unites internalist and externalist elements. Especially relevant to Chapters 1, 2, and 7–11.
- , *Perceiving God: The Epistemology of Religious Experience*, Ithaca, 1991. A major study in religious epistemology, also containing much general epistemology, particularly in the theory of perception. Highly relevant to Chapters 1, 2, and 12.
- , *Beyond "Justification,"* Ithaca, 2005. An account of the various elements a sound epistemology should take into account. Relevant especially to Chapters 1, 2, and 11.
- Aristotle, *Posterior Analytics*, trans. by G.R.G. Mure, Oxford, 1928. A major text of immense influence. Especially relevant to Chapter 9.
- Armstrong, D.M., *Belief, Truth and Knowledge*, Cambridge, 1973. A wide-ranging epistemological essay influential in the reliabilist tradition. Especially pertinent to Chapters 1–3 and 9–11.
- Audi, Robert, *The Structure of Justification*, Cambridge and New York, 1993. A collection of papers covering (and often clarifying) most of the topics in Chapters 4, 6–8, and 10–14.
- , *The Architecture of Reason*, Oxford and New York, 2001. A comprehensive theory of rationality, practical as well as theoretical; explicates in detail the nature of justification by comparison with rationality; offers an account of the role of experience in grounding both; and applies some of its epistemological results to the foundations of ethics. Particularly relevant to Chapters 1, 4, 10, and 12.
- , *The Good in the Right: A Theory of Intuition and Intrinsic Value*, Princeton, 2004. A statement and defense of an intuitionist ethics, with much discussion of intuition, self-evidence, and moral knowledge. Relevant mainly to Chapters 5, 6, 10, and 12.
- Aune, Bruce, *Knowledge of the External World*, London, 1991. A historically oriented, integrated treatment of empirical knowledge. Bears particularly on Chapters 1, 2, 10, and 11.
- Austin, J.L., *Sense and Sensibilia*, Oxford, 1979 (reprint). A study of perception particularly pertinent to the appraisal of sense-datum theory and of relevance to understanding foundationalism. Especially relevant to Chapters 1, 2, and 9.
- Ayer, A.J., *The Problem of Knowledge*, Harmondsworth, 1956. An epistemological survey in the empiricist tradition. Especially relevant to Chapters 1, 2, 8, 10, and 11.

- Bergmann, Michael, *Justification without Awareness*, Oxford and New York, 2006. This book defends an externalist theory of justification with intensive scrutiny of internalist awareness conditions for justification and a development of a proper-function view of justification and knowledge. Highly relevant to Chapters 11 and 12.
- Berkeley, George, *A Treatise Concerning the Principles of Human Knowledge*, New York, 1929 (originally published in 1710). Perhaps the most important statement of phenomenalism, and of an idealist world view in general, in the empiricist tradition. Bears particularly on Chapters 1–3, 8, 9, and 10.
- Bernecker, Sven, *Memory: A Philosophical Study*, Oxford and New York, 2009. A detailed presentation of a theory of memory and its role in knowledge and justification. Highly relevant to Chapters 3, 10, and 11.
- Boghossian, Paul and Peacocke, Christopher (eds.), *New Essays on the A Priori*, Oxford and New York, 2000. A collection of professional papers on epistemological, metaphysical, and logical and semantic aspects of the a priori, with an introduction by the editors and separate chapters by each. Pertinent above all to Chapters 4–5.
- BonJour, Laurence, *The Structure of Empirical Knowledge*, Cambridge, MA, 1985. A detailed and widely discussed statement and defense of coherentism regarding empirical knowledge. Especially pertinent to Chapters 1, 2, 8, 9, and 10.
- , *In Defense of Pure Reason*, Cambridge and New York, 1998. An account of a priori justification and knowledge, with detailed critical discussions of Kant's position on them and of Hume on induction. Highly pertinent to Chapters 5, 6, and 10.
- , *Epistemology: Classical Problems and Contemporary Responses*, 2nd edn, Lanham, MD, 2009. An introductory book covering most (though not all) of the major topics treated in this book and defending a different though overlapping view of perceptual justification.
- Brandt, R.B. and Nagel, Ernest (eds.), *Meaning and Knowledge*, New York, 1965. A comprehensive, historically informed assembly of classical and recent readings in epistemology. Bears on all the chapters.
- Butchvarov, Panayot, *The Concept of Knowledge*, Evanston, IL, 1970. A metaphysically oriented, foundationalist inquiry into the nature of knowledge. Especially pertinent to Chapters 1, 2, 4, and 8–10.
- , *Skepticism in Ethics*, Bloomington, IN, 1989. A detailed study of the possibility and possible extent of justification and knowledge in ethics. Particularly relevant to Chapters 12–14.
- , *Skepticism and the External World*, Oxford and New York, 1998. An account of skepticism and its implications for a realist metaphysics, with a distinctive theory of the role of decision in determining the extent of our knowledge. Most closely pertinent to Chapters 11, 13, and 14.
- Chisholm, R.M., *Theory of Knowledge*, 2nd and 3rd edns, Englewood Cliffs, NJ, 1977 and 1989. A rigorous treatment of many basic topics in the field, and an introduction to Chisholm's own epistemological views. Relevant to each chapter.
- , *The Foundations of Knowing*, Minneapolis, 1982. A collection of essays, on topics relevant to each chapter, by one of the major epistemologists of the twentieth century.
- Chisholm, R.M. and Swartz, Robert (eds.), *Empirical Knowledge*, Englewood Cliffs, NJ, 1973. A large collection of influential epistemological writings. Especially pertinent to Chapters 1, 2, 3, and 7–10.

- Coady, C.A.J., *Testimony*, Oxford, 1992. A comprehensive, historically wide-ranging treatment of testimony that links it at many points with the philosophy of language and the views of Wittgenstein. Pertinent above all to Chapter 7.
- Craig, Edward, *Knowledge and the State of Nature*, Oxford, 1990. A sophisticated but lucid and concrete treatment of the problem of analyzing knowledge, with considerable attention to skepticism and far more on testimony than most general books on epistemology. Particularly relevant to Chapters 7, 11, and 13.
- Crumley, Jack S. II, *An Introduction to Epistemology*, 2nd edn, Petrboro, Ontario, 2009. A wide-ranging but concise introduction to both the problems and literature of the field. Pertinent to nearly every chapter.
- (ed.), *Readings in Epistemology*, Mountainview, CA, 1999. A comprehensive collection of mainly contemporary readings in the field, with introductions to each section. Bears on each chapter.
- Dancy, Jonathan, Sosa, Ernest, and Matthias Steup (eds.), *A Companion to Epistemology*, 2nd edn, Oxford, 1992. A comprehensive reference work of short, dictionary- and encyclopedia-style articles. Highly pertinent to every chapter.
- DePaul, Michael R. and Ramsey, William (eds.), *Rethinking Intuition: The Psychology of Intuition and its Role in Philosophical Inquiry*, Lanham, MD, 1998. A collection of papers by philosophers and psychologists on the nature of intuition, with pertinent discussion of recent psychological studies of reasoning. Especially relevant to Chapters 4, 5, and 12 but also useful in understanding philosophical method in general.
- DeRose, Keith, *The Case for Contextualism*, Oxford and New York, 2009. A detailed account of contextualism with attention to many variants and a partial theory of the importance of the view for skepticism. Especially relevant to Chapters 11, 13, and 14.
- DeRose, Keith and Warfield, Ted A. (eds.), *Skepticism: A Contemporary Reader*, Oxford and New York, 1999. A collection of recent writings that concern skepticism particularly in the context of questions about contextual variables, meaning, self-knowledge, and transmission of justification. Highly pertinent to Chapters 8, 11, 13, and 14.
- Descartes, René, *Meditations on First Philosophy*, in vol. 2 of *The Philosophical Writings of Descartes*, trans. by John Cottingham, Robert Stoothof, and Dugald Murdoch, Cambridge, 1984 (originally published in 1641). One of the greatest and most influential works in modern epistemology, and a powerful statement of both rationalism and foundationalism. Bears on every chapter.
- Dretske, Fred I., *Seeing and Knowing*, London, 1969. An intensive study of perception and its relation to empirical knowledge, with an influential defense of the notion of non-epistemic seeing. Especially relevant to Chapters 1, 2, and 11–14.
- , *Knowledge and the Flow of Information*, Cambridge, MA, 1981. A major statement of a reliabilist, information-theoretic account of knowledge. Particularly pertinent to Chapters 1, 8, and 10.
- Everitt, Nicholas and Fisher, Alec, *Modern Epistemology*, New York, 1995. A high-level introductory textbook with special attention to mathematical and logical knowledge and detailed discussions of Quine (a major influence on the authors) and Rorty. Pertinent to most chapters, but especially Chapters 5, 6, and 8.
- Fales, Evan, *A Defense of the Given*, Lanham, MD, 1966. A defense of the view that justification is conferred by what is given in immediate experience, with detailed applications to the foundationalism conceived as a plausible theory that makes use of the given. Pertinent to Chapters 1, 2, and 9.

- Feldman, Richard, *Epistemology*, Upper Saddle River, NJ, 2003. An Introductory treatment relevant mainly to Parts One and Three and especially to Chapters 10 and 11.
- Fogelin, Robert, *Pyrrhonian Reflections on Knowledge and Justification*, Oxford and New York, 1994. An intensive and sympathetic inquiry into skepticism, from a distinctively Pyrrhonian point of view. Relevant above all to Chapters 13 and 14.
- Foley, Richard, *The Theory of Epistemic Rationality*, Cambridge, MA, 1987. A detailed development of a subjectivistic theory of justification. Pertinent mainly to Parts Two and Three.
- , *Working without a Net*, Cambridge, MA, 1993. A development and generalization of his theory of epistemic rationality. Especially relevant to Chapters 11–14.
- Fumerton, Richard A., *Metaphysical and Epistemological Problems of Perception*, Lincoln, NE, 1985. A study of both the nature of the objects of perception and the way perception yields justification and knowledge. Especially pertinent to Chapters 1, 2, and 10.
- , *Metaepistemology and Skepticism*, Lanham, MD, 1996. A systematic treatment of skepticism with special reference to the controversies between internalism and externalism and foundationalism and coherentism. Highly relevant to Chapters 11 and 12–14.
- Ginet, Carl, *Knowledge, Perception, and Memory*, Dordrecht and Boston, 1975. A rigorous high-level study of these three notions in relation to each other. Particularly relevant to Chapters 1–3 and 11–14.
- Goldman, Alan H., *Empirical Knowledge*, Berkeley, CA, 1988. A detailed presentation of a theory of knowledge and evidence with special emphasis on inference to the best explanation. Pertinent above all to Chapters 1 and 9–14.
- Goldman, Alvin I., *Epistemology and Cognition*, Cambridge, MA, 1986. A presentation of reliabilism for justification and knowledge, with much discussion of related developments in cognitive psychology. Particularly relevant to Chapters 1, 3, 8, 10, and 11.
- , *Knowledge in a Social World*, Oxford and New York, 1999. A wide-ranging treatment of epistemology in the context of social elements in its development, with much discussion of the nature and structure of knowledge in such domains as science, law, and education. Bears on many chapters, particularly 7 and 12.
- , *Pathways of Knowledge*, Oxford and New York, 2002. A concise treatment of a number of the topics treated in this book.
- Greco, John, *Putting Skeptics in Their Place*, Cambridge and New York, 2000. An examination of the structure of skeptical arguments which contends that taking account of them requires an externalist epistemology and that a version of virtue epistemology deals best with knowledge and skepticism. Particularly pertinent to Chapters 1, 10, 13, and 14.
- Greco, John and Sosa, Ernest, *The Blackwell Guide to Epistemology*, Oxford, 1999. A book of commissioned papers aimed at providing high-level examinations of the current status of major epistemological issues. Pertinent to every chapter.
- Haack, Susan, *Evidence and Inquiry*, Oxford, 1993. An attempt to combine the best elements in foundationalism and coherentism, with special attention to the connections between epistemology and philosophy of science. Bears especially on Chapters 8, 9, and 12.
- Hamlyn, David, *Understanding Perception*, Aldershot, 1996. A monograph on

- perception and its relation to belief, learning, consciousness, and other notions important in this book. Highly relevant to Chapters 1, 2, and 4.
- Harman, Gilbert, *Thought*, Princeton, NJ, 1975. A broad, influential study of the nature of knowledge, with much attention to defeasibility conditions, inferential transmission of knowledge, and the role of coherence in justification. Bears especially on Chapters 1, 3, and 8–11.
- , *Change in View*, Cambridge, MA, 1986. A detailed study of evidence and justification as bearing on changing, as opposed to simply holding, a position, with special emphasis on the conditions for coherence in one's overall view. Especially relevant to Chapters 8–11.
- , *Reasoning, Meaning, and Mind*, Oxford and New York, 1999. A collection of papers presenting large-scale conceptions of the a priori, reasoning, and knowledge; it helpfully contrasts with the perspectives in Chapters 5, 8, and 9.
- Hawthorne, John, *Knowledge and Lotteries*, Oxford and New York, 2004. A study of the lottery problem and the conditions for the transmission of knowledge, with direct bearing on Chapter 8 and relevance to Chapters 10–14.
- Heatherington, Stephen Cade, *Knowledge Puzzles*, Boulder, CO, 1996. A brief but comprehensive introduction to epistemology with emphasis on the puzzles generated by many of its central questions. One or another section bears on any given chapter.
- Heil, John, *Perception and Cognition*, Berkeley, CA, 1983. A study of perception in relation to belief, with discussions linking epistemology to philosophy of mind. Especially pertinent to Chapters 1, 2, 4, and 10.
- Hintikka, Jaakko, *Knowledge and Belief*, Ithaca, NY, 1962. A rigorous, advanced study in epistemic logic. Particularly relevant to Chapters 8 and 10.
- Huemer, Michael, *Skepticism and the Veil of Perception*, Lanham, MD, 2001. A critical, historically informed analysis of skepticism about perceptual knowledge, with a detailed response in the foundationalist, realist tradition. Highly pertinent to Chapters 1, 2, 9, and 10.
- (ed.), *Epistemology: Contemporary Readings*, London, 2002. A large, wide-ranging collection of classical and contemporary papers and book selections. It is designed to complement this book and has at least several readings pertinent to every chapter, as well as introductions to each major section by Huemer and useful study questions for each reading.
- Hume, David, *A Treatise of Human Nature*, Oxford, 1888 (originally published in 1739). One of the greatest and most influential works in modern philosophy, and a powerful statement of empiricism. It bears on every chapter.
- , *An Enquiry Concerning Human Understanding*, Indianapolis, 1977 (first published in 1748). A major work in modern epistemology, particularly for the topics of causation, induction, and skepticism. Especially relevant to Chapters 1, 2, 10, and 11.
- Kant, Immanuel, *Prolegomena to Any Future Metaphysics*, trans. by Lewis White Beck, New York, 1950 (originally published in 1783). A short presentation of the ideas in Kant's monumental *Critique of Pure Reason*, one of the greatest texts in modern philosophy. The *Prolegomena* bears particularly on Chapters 1, 5, 6, 10, and 13.
- Kitcher, Philip, *The Nature of Mathematical Knowledge*, Oxford, 1984. A wide-ranging, historically informed treatment of mathematical knowledge, with a sustained



- critique of various apriorist conceptions of it and an account of a “defensible empiricism.” Particularly relevant to Chapters 4 and 5.
- Klein, Peter, *Certainty: A Refutation of Scepticism*, Minneapolis, 1981. A rigorous treatment of the nature of knowledge, certainty, and skepticism. Especially pertinent to Chapters 8 and 10–14.
- Kvanvig, Jonathan L., *The Intellectual Virtues and the Life of the Mind*, Lanham, MD, 1992. A detailed statement and critical appraisal of a virtue epistemology, relevant mainly to Chapters 1, 10, and 11.
- (ed.), *Warrant in Contemporary Epistemology: Essays in Honor of Plantinga’s Epistemology*, Lanham, MD, 1996. A collection of commissioned papers on the warrant approach developed by Plantinga and related topics. Bears particularly on Chapters 10 and 11.
- Kyburg, Henry E., Jr., *Epistemology and Inference*, Minneapolis, 1983. A collection of advanced papers, of which several are especially pertinent to Chapters 10–12.
- Lackey, Jennifer, *Learning from Words: Testimony as a Source of Knowledge*, Oxford and New York, 2008. A comprehensive theory of the nature and epistemological significance of testimony. Particularly pertinent to Chapters 7 and 11.
- Lackey, Jennifer and Sosa, Ernest (eds.), *The Epistemology of Testimony*, Oxford and New York, 2006. Contains contemporary essays squarely relevant to Chapter 7 and bearing on Chapters 11 and 12.
- Lehrer, Keith, *Knowledge*, Oxford, 1974. A critique of foundationalism and development of the author’s own coherentist theory of knowledge, with a detailed, non-causal account of justification. Particularly relevant to Chapters 1, 4, and 9–11.
- , *Theory of Knowledge*, Boulder, CO, 1990. A successor to his *Knowledge* that presents a revised theory and remains relevant to Chapters 1, 4, and 9–11.
- Lewis, C.I., *An Analysis of Knowledge and Valuation*, La Salle, IL, 1946. A systematic study, in the foundationalist tradition, of many major epistemological and metaphysical questions. Especially relevant to Chapters 1–6 and 7–11.
- Locke, Don, *Memory*, New York, 1971. A concise but wide-ranging treatment of memory and memorial knowledge. Bears mainly on Chapter 3.
- Locke, John, *An Essay Concerning Human Understanding*, New York, 1928 (originally published in 1689). A major text in modern epistemology, and a powerful statement of a commonsense empiricism. Bears on each chapter.
- Luper-Foy, Stephen, *The Possibility of Knowledge: Nozick and His Critics*, Lanham, MD, 1987. A set of papers discussing above all Nozick’s tracking account of knowledge and its bearing on skepticism. Relevant especially to Chapters 8, 10, and 11.
- Lycan, William, *Judgment and Justification*, New York, 1988. An account of justification and of aspects of knowledge, with emphasis on their connection with the notion of explanation and with selected topics in the philosophy of mind. Especially relevant to Chapters 1, 2, and 8–11.
- Lyons, Jack C., *Perception and Basic Beliefs: Zombies, Modules, and the Problem of the External World*, Oxford and New York, 2009. A monographic, psychologically informed study of perception defending a reliabilistic account of its nature and role in grounding knowledge. Relevant especially to Chapters 1, 2, 8, and 11–14.
- McGrew, Timothy J., *The Foundations of Knowledge*, Lanham, MD, 1995. A systematic defense of foundationalism (strong as opposed to moderate); relevant mainly to Chapters 1, 5, 8, and 9.
- Malcolm, Norman, *Knowledge and Certainty*, Ithaca, NY, 1975. A collection of

- epistemological essays by a leading proponent of the philosophy of Wittgenstein. Especially relevant to Chapters 1–4, 10, and 12.
- Mill, John Stuart, *A System of Logic*, London, 1843. The leading nineteenth-century statement of radical empiricism. Especially relevant to Chapters 1, 2, and 6.
- Montmarquet, James A., *Epistemic Virtue and Doxastic Responsibility*, Lanham, MD, 1993. A statement and appraisal of a virtue epistemology with special emphasis on the relation between justification and intellectual responsibility. Especially relevant to Chapters 1, 10, and 11.
- Morton, Adam, *A Guide through the Theory of Knowledge*, 3rd edn, Oxford, 2002. A brief but moderately comprehensive introduction, with more on probability and induction than many short books of the same range. Pertinent to nearly every chapter.
- Moser, Paul K., *Knowledge and Evidence*, Cambridge and New York, 1993. A rigorous and detailed presentation of an account of knowledge and the nature and scope of its supporting grounds. Especially pertinent to Chapters 1, 2, and 9–11.
- , *Philosophy after Objectivity*, Oxford and New York, 1993. A rigorous treatment of the relation of realism to skepticism and of both to semantic questions about meaning and epistemological questions about evidence, reasons, and justification. Especially pertinent to Chapters 5, 6, and 10–12.
- (ed.), *The Oxford Handbook of Epistemology*, Oxford and New York, 2002. A collection of wide-ranging papers covering central issues and developments in most of the main the areas treated in this book.
- Moser, Paul K. and van der Nat, Arnold (eds.), *Human Knowledge*, 2nd edn, New York, 1995. An extensive set of classical and contemporary readings, with editors' introductions to each section. Pertinent to every chapter.
- Moser, Paul K., Mulder, Dwayne H., and Trout, J.D. (eds.), *The Theory of Knowledge*, Oxford, 1998. A brief introduction covering many of the topics in this book and especially pertinent to understanding the relation between epistemology and philosophy of science. Bears on most chapters.
- Nathan, N.M.L., *The Price of Doubt*, London, 2000. A systematic treatment of skepticism, with much examination of the notion of evidence and rational belief. Pertinent especially to Chapters 10, 11, 13, and 14.
- Neta, Ram, and Pritchard, Duncan (eds.), *Arguing about Knowledge*, London and New York, 2009. A comprehensive anthology covering most of the topics in this book and with editorial introductions to each of its many sections.
- Niiniluoto, Ilkka, Matti Sintonen, and Jan Wolenski (eds.), *Handbook of Epistemology*, Dordrecht, 2004. A large collection of essays on the major areas of epistemology by. Pertinent to every chapter.
- Nozick, Robert, *Philosophical Explanations*, Cambridge, MA, 1981. Contains a book-length epistemology section bearing on Chapters 9–11, 13, and 14.
- Pappas, George S. and Swain, Marshall, *Essays on Knowledge and Justification*, Ithaca, NY, 1978. A collection of papers, many influential, mainly from professional journals. Especially relevant to Chapters 1 and 10–14.
- Peacocke, Christopher, *Being Known*, Oxford, 1999. A philosophical treatment of knowledge, especially self-knowledge, in relation to central issues in philosophy of language and metaphysics, particularly intentionality, representation, and truth. Highly pertinent to Chapters 3, 4, 10, and 11.
- Plantinga, Alvin, *Warrant and Proper Function*, Oxford and New York, 1993. A systematic development of a theory of warrant as to what renders true belief

- knowledge and of key epistemic concepts related to warrant. Particularly relevant to Chapters 1, 5, 6, and 10–14.
- , *Warrant: The Current Debate*, Oxford and New York, 1993. A critical study of many major epistemological theories. Highly relevant to Chapters 8–14.
- Plato, *Theaetetus*, trans. by F.M. Cornford, Cambridge, 1934. A major text in the history of epistemology. Especially relevant to Chapter 10.
- Pojman, Louis P. (ed.), *Knowledge: Classic and Contemporary Readings*, Belmont, CA, 1993. A large collection, with introductions to each section. Relevant to every chapter.
- , *What Can We Know?*, 2nd edn, Belmont, CA, 2001. A wide-ranging, highly readable, historically informed introduction with attention both to basic issues and to classical and contemporary literature. Pertinent to every chapter.
- Pollock, John L., *Knowledge and Justification*, Princeton, 1974. A systematic, wide-ranging treatment of these notions which develops a foundationalist account of both. It bears on all the chapters.
- Price, H.H., *Perception*, Oxford, 1932. A detailed major study of the topic, with important discussions of sense-data. Especially relevant to Chapters 1, 2, 10, and 11.
- Pritchard, Duncan, *Epistemic Luck*, Oxford and New York, 2005. An exploration of how knowledge must be conceived given an understanding of the sense in which it may not be due to luck and of how the resulting conception bears on skepticism. Particularly pertinent to Chapters, 10, 11, and 13.
- Quine, W.V., *The Pursuit of Truth*, revised edn, Cambridge, MA, 1992. A short statement of many of Quine's epistemological views, with valuable discussions connecting epistemology with the philosophy of science. The chapters devoted to evidence, meaning, perception, and truth make it especially pertinent to Chapters 1, 5, 6, 10, and 12.
- Quine, W.V. and Ullian, Joseph, *The Web of Belief*, 2nd edn, New York, 1978. An introductory treatment of the development of commonsense and scientific knowledge. Bears particularly on Chapters 1, 2, and 8–10.
- Quinton, Anthony, *The Nature of Things*, London, 1973. A survey of problems in metaphysics and epistemology, with much discussion of the foundationalism–coherentism controversy. Especially relevant to Chapters 1, 9, and 10.
- Reid, Thomas, *Essays on the Intellectual Powers of Man*, London, 1869 (originally published in 1785). An important critique of other modern philosophers, rationalist and empiricist, and an original development of a commonsense epistemology. Especially relevant to Chapters 1–7 and 10–14.
- Rescher, Nicholas, *Skepticism*, Oxford, 1980. A detailed critical survey of many of the arguments for skepticism. Especially relevant to Chapters 10–14.
- Robinson, Howard, *Perception*, London, 1994. A historically detailed defense of the sense-datum theory in the philosophy of perception. Bears on Chapters 1–4.
- Rorty, Richard, *Philosophy and the Mirror of Nature*, Princeton, 1979. A widely discussed critique of the foundationalist tradition in epistemology (and indeed of aspects of epistemology itself), with special attention to foundationalism as deriving from Descartes. Pertinent especially to Chapters 1 and 9.
- Russell, Bertrand, *The Problems of Philosophy*, London, 1912. An influential introductory and highly readable survey of epistemology and metaphysics. Particularly pertinent to Chapters 1, 4, and 10–14.
- , *An Inquiry into Meaning and Truth*, London, 1940. An articulation of an overall

- epistemological position by one of the twentieth century's major epistemologists. Especially relevant to Chapters 1, 2, and 8–11.
- Sellars, Wilfrid, *Science, Perception, and Reality*, London, 1963. A collection of advanced epistemological papers by a major contemporary philosopher and leading coherentist. Especially pertinent to Chapters 1–4 and 9–12.
- Shope, Robert, *The Analysis of Knowing*, Princeton, 1983. A highly detailed study of major kinds of analysis of knowledge. Especially relevant to Chapters 10 and 11.
- Sinnott-Armstrong, Walter, *Moral Skepticisms*, Oxford and New York, 2006. A development of a contrastivist epistemology, with detailed analysis of its bearing on the possibility of moral knowledge and an account of the scope and varieties of plausible skepticisms. Pertinent especially to Chapters 8, 10, and 12–14.
- Sosa, Ernest, *Knowledge in Perspective*, Cambridge and New York, 1991. A wide-ranging collection of papers intensively exploring most of the topics treated in Chapters 1, 2, 3, 10, and 11, and introducing a sophisticated and influential version of virtue epistemology.
- , *Virtue Epistemology: Apt Belief and Reflective Knowledge*, Oxford and New York, 2007. A detailed statement of Sosa's version of virtue epistemology, with special relevance to Chapters 1, 8, and 10–12.
- Sosa, Ernest and Kim, Jaegwon (eds.), *Epistemology: An Anthology*, Oxford and New York, 2000. A wide-ranging set of mainly recent writings on justification, knowledge, and skepticism. Highly pertinent to Chapters 1, 6, 8, and 10.
- Steup, Matthias, *An Introduction to Contemporary Epistemology*, Englewood Cliffs, NJ, 1996. A comprehensive and systematic introduction to epistemology, with close attention both to theoretical and conceptual issues and to major statements in the literature.
- , *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, Oxford and New York, 2001. A collection of recent papers focusing largely on the relation between belief and will and on how that relation bears on knowledge and justification in general and on virtue epistemology in particular. Particularly relevant to Chapters 10 and 11.
- , *Contemporary Debates in Epistemology*, Oxford, 2005. A volume in which prominent contemporary epistemologists debate major issues, including many central in this book. Especially relevant to Chapters 8–14.
- Stroud, Barry, *The Significance of Philosophical Scepticism*, Oxford, 1984. An intensive study of skepticism, particularly as set out by Descartes. Particularly relevant to Chapters 1, 2, 10, 12, and 13.
- Swain, Marshall, *Reasons and Knowledge*, Ithaca, NY, 1981. A rigorous statement of an account of knowledge that unites justificationist and reliabilist elements. Especially relevant to Chapters 1, 2, 8, 10, and 11.
- Swinburne, Richard, *Epistemic Justification*, Oxford and New York, 2001. A wide-ranging treatise with much on the contrast between internalism and externalism; it clarifies the importance of probability for justification and explores the relation between the genesis of beliefs and their epistemic standing. Particularly pertinent to Chapters 8 and 9–11.
- Unger, Peter, *Ignorance*, Oxford, 1975. A full-scale case for skepticism. Bears particularly on Chapters 1, 2, 10, 12, and 13.
- Vinci, Thomas, *Cartesian Truth*, Oxford and New York, 1998. An interpretation of Descartes's epistemology that rigorously defends some main Cartesian ideas. Bears especially on Chapters 1, 2, 5, 8, 10, and 11.

- Will, Frederick L., *Induction and Justification*, Ithaca, NY, 1974. A critical study of a foundationalist approach to justification, with much attention to the problem of induction. Especially relevant to Chapters 9–14.
- Williams, Michael, *Unnatural Doubts*, Oxford, 1991. A detailed diagnostic account of skepticism and some major issues concerning the possibility of foundations of knowledge. Pertinent especially to Chapters 1, 2, and 9–12.
- , *Problems of Knowledge*, Oxford and New York, 2001. A broad treatment of many of the problems of epistemology, with particular attention to skepticism. Pertinent to nearly every chapter.
- Williamson, Timothy, *Knowledge and its Limits*, Oxford and New York, 2000. A rigorous defense of the idea that knowledge is a state of mind and should be taken as a basis for understanding other notions important in epistemology, such as evidence, rather than analyzed in terms of allegedly more basic notions. Especially pertinent to Chapters 1, 2, 8–11, and 13.
- Wittgenstein, Ludwig, *Philosophical Investigations*, trans. by G.E.M. Anscombe, Oxford, 1953. Wittgenstein's most wide-ranging and influential work, spanning topics in epistemology, metaphysics, and other philosophical areas. Its remarks and examples bear on every chapter.
- , *On Certainty*, trans. by Dennis Paul and G.E.M. Anscombe, Oxford, 1969. An important discussion of the topic by one of the major twentieth-century philosophers. Especially relevant to Chapters 9–14.
- Zagzebski, Linda, *Virtues of the Mind*, New York, 1996. A statement of virtue epistemology distinctive for its background account of virtue ethics as the counterpart on which virtue epistemology should be based. Pertinent above all to Chapters 10 and 11.

# Index

Page numbers in **bold** show where a term is defined.

- a posteriori (empirical) propositions **114**
- a priori justification 141–4
- a priori knowledge 141–4
- a priori proposition 111, 112–14, 116, 124 (n. 11), 125 (n. 14), 140, **144**, 225, 256
- abduction *see* inference to the best explanation
- abstract entities 116, 120, 256, 292 (n. 4)
- acceptance 326–7; *see also* belief
- accessibility 87, 189, 207, 316; internal 272–8; perceptual 322; public 304, 306
- accidentality 31, 64, 270, 272, 283, 284; *see also* chance, luck
- adverbial theories: applied to hallucination 48–49; applied to memory 70–71, 74; of perception 47; of sensory experience 48–50; *see also* phenomenalism
- Alston, William P. 59 (n. 1), 236 (n. 1), 266 (n. 134), 292 (n.5), 294 (nn. 15, 17), 331 (n. 25), 376 (n. 8)
- Amico, Robert 355 (n. 13)
- analysis 54, 107, 110–11, 246, 247, 291, 295
- analytic proposition 107–8, 112, 117, 130; as independent of experience 113
- Anderson, Elizabeth 329 (n. 9)
- anti-realism 89, 124 (n. 10), 170 (n. 13)
- appearance state 240 (n. 20)
- Aquinas, St Thomas xi, 123 (n. 6), 125 (n. 13), 323 (n. 9), 343, 355 (n. 9)
- argument from analogy 241
- Aristotle 121, 215, 238 (n. 7), 277, 278, 294 (n. 15)
- Armstrong, D.M. 34 (n. 10), 59 (n. 4), 265 (n. 9)
- attesting 151
- Austin, J.L. 201 (n. 9)
- Averill, Edward Wilson 124 (n. 10)
- Ayer, A.J. 201 (n. 9), 292 (n. 6)
- Barnes, Winston H.F. 60 (n. 8)
- basic belief **155**, 165
- basic knowledge 155, 160
- basic source 6–7, **78**, 138, **154**, 155, 167, 184, 191, 216, 234, 320, 381
- basing relation 68, 106, 142, 147 (n. 6), 154, 168 (n. 6), 180, 181, 199, 200 (nn. 5, 7), 208, 209, 219, 220; direct vs. indirect 143
- begging the question 339, 346, 348, 349, 354 (n. 5)
- belief: and acceptance 326–7; about perceptibles **16** (*see also* belief, perceptual); a priori 8, 139–141, **144**; as dispositional 85–6, 231; basic *see* basic belief; episodically vs. structurally inferential beliefs 180; dispositional vs. disposition to believe 34 (n. 8), 153, 231, 314; direct 143, **182**; for a reason 178, 200 (n. 6); foundational (*see also* basic belief) 216, 233, 235, 239 (n. 16); genesis vs. justification 118; indirect 181, 210; inductive 340, 361; infallible 92, 94, 345, 346, 347; inferential 178, 180; inferential development of 181, 185 (*see also* belief, inferential); innate 65, 79 (n. 5); introspective 214–15; memorial 8, 364; non-inferential 59 (n.4), 124 (n. 11), 152, 160, 178, 182, 198; objectual 21, 22, 27, 32 (n. 2); perceptual **16**; propositional 32 (nn. 2, 3, 4), 91, 105; reasoned 178–80; retained 63–4; second-order 229–30, 367; self-grounding **346**, 361; testimony-based, 151–2, 154–55
- Bender, John W. 239 (n. 15)
- Berkeley, George 52, 53, 60 (n. 10)
- Bernecker, Sven 81 (n. 12), 240 (n. 19)
- Blanshard, Brand 294 (n.18)
- blind sight 34 (n. 9), 57, 58
- BonJour, Laurence 59 (n. 3), 122 (n. 1), 126 (n. 24), 238 (n. 11), 239 (n. 15)
- brain in a vat 242, 353 (n. 1), 355 (n. 8)
- Brandt, Richard B. 330 (n. 20)
- Brueckner, Anthony 355 (n. 8)
- Burge, Tyler 172 (n. 17)
- Butchvarov, Panayot 354 (n. 7)

- Calvin, John 320  
 Canary, Catherine 203 (n. 22)  
 Castañeda, Hector-Neri 238 (n. 8), 240 (n. 19)  
 Casullo, Albert 146 (n. 4)  
 categorical imperative 315  
 causal dependence 55, 260  
 causal theory of perception *see* perception, causal theory of  
 certainty 236, 265 (n. 8), 335, 347–9; certainty principle 347; epistemic 251, 348; propositional 251; psychological 250  
 chance 191, 249, 259–60; *see also* accidentality, luck  
 Chisholm, R.M. 59 (n. 6), 60 (nn. 8, 9), 292 (n. 6), 355 (n. 13)  
 circularity problem 370–1, 376 (n. 8)  
 clairvoyance 270–2  
 Coady, C.A.J. 167 (n. 1), 171 (n. 17), 172 (n. 22)  
 Code, Lorraine 329 (n. 8)  
 Cohen, Stewart 266 (n. 15)  
 coherence 216–34; with experience 235, 240 (n. 20); *see also* coherentism  
 coherence theory of concepts 226–7  
 coherentism 216–32; holistic 218–21; isolation problem for 222–3; linear 217; *see also* coherence  
 common sense 40, 54, 134, 214, 325, 346, 349; negative defense of 358; positive defense of 358  
 conative attitudes 33 (n. 4)  
 conception 19–21, 27  
 Conee, Earl 292 (n. 5)  
 consciousness 84–7, 91–3, 170 (n. 14), 273; *of vs. that* 285  
 consequentialism 317  
 conservatism, epistemic 325–6  
 conservatism, phenomenal 147 (n. 7), 326  
 content: belief, 25, 32 (n. 2), 33 (n. 7), 34 (n. 10), 35 (n. 14); perceptual 25, 27–8; property, 27–8; propositional 28  
 contextualism 261–3, 266 (n. 15), 356 (n. 14)  
 contingent proposition 106  
 contrast class 356 *see also* relativism, relevant alternative  
 contrastivism 263, 266 (n. 18)  
 conventionalism 130–3  
 correspondence theory of truth *see* truth, correspondence, theory of  
 courage, intellectual 373–4  
 credibility, testimonial 151–4, 156  
 credulity 152  
 criterion, problem of the *see* problem of the criterion  
 Crumley, Jack S., II 329 (nn. 8, 9)  
 Davidson, Donald 123 (n. 7), 171 (n. 17), 240 (n. 19), 294 (n. 18)  
 de Almeida, Claudio 202 (n. 13), 237 (n. 4)  
*de dicto* 32 (n. 3)  
 deducibility 313  
 deductive reasoning 186–7; and proof 299; and scientific knowledge 302  
 deductive transmission 188, 193–7  
 defeasibility 105, 189, 201 (n. 13); *see also* justification, indefeasible  
 definition, truth by 130–2; *see also* meaning  
 dependence *see* causal dependence; discriminative dependence; epistemic dependence; evidential dependence; functional dependence; genetic dependence; inferential dependence; justificational dependence; operational dependence; presuppositional dependence; psychological dependence  
*de re* 32 (n. 3)  
 DeRose, Keith 266 (nn. 15, 18)  
 Descartes, René 42, 79 (n. 3), 92, 99, 216, 239 (n. 18), 289, 342, 343, 345, 348, 353 (n. 1), 355 (n. 11)  
 design 265 (n. 12), 267, 319  
 Dewey, John 294 (n. 19)  
 Dickinson, Emily 42–3  
 disagreement, rational 311, 313, 371–4  
 discriminative dependence 35, 57  
 disjunctivism 59 (n. 2)  
 disposition to believe *see* belief, dispositional vs. disposition to believe  
 dispositional belief *see* belief, dispositional vs. disposition to believe  
 dispositional properties 85–6  
 disquotation principle 294 (n. 16)  
 dogmatism 236, 364, 374  
 doxastic attitudes 33 (n. 5)  
 doxastic voluntarism 207–8, 236 (n. 1)  
 dream argument 353 (n. 1)  
 Dretske, Fred 34 (n. 11), 35 (n. 14), 203 (n. 22), 240 (n. 19), 265 (n. 7)  
 Duhem, Pierre 328 (n. 2)  
 Duhem–Quine thesis 328 (n. 2)  
 egocentric predicament 342–3  
 Einstein, Albert 118  
 empirical *see* a posteriori proposition  
 empiricism 116, 120  
 entailment 115, 125 (n. 13), 126 (n. 20)  
 envatment problem 242  
 epistemic conservatism *see* conservatism, epistemic  
 epistemic authority 362–3  
 epistemic chain 211  
 epistemic dependence 228–9; negative 228–9; positive 228–9; *see also* causal dependence; presuppositional dependence; psychological dependence  
 epistemic grounding *see* grounding, epistemic  
 epistemic immediacy *see* immediacy, epistemic  
 epistemic indirectness *see* indirectness  
 epistemic parity 372–3  
 epistemic particularism 201 (n. 10)



- epistemic possibility 360, 375 (n. 3)  
 epistemic power 143, 278  
 epistemic principle 30  
 epistemic priority 306  
 epistemic regress argument 215  
 epistemic regress problem 210  
 epistemic virtue 277–9  
 epistemology, naturalistic *see* naturalistic epistemology  
 essential truth *see* truth, essential  
 Euclidean geometry 194  
 evidence of the senses 28  
 evidential dependence 112  
 evidentialism 320  
 excluded middle, principle of the 120, 121  
 excusability 76, 316,  
 experientialism 320–5  
 explanation 111, 221; mutual 225–6; reciprocal 226; *see also* inference to the best explanation  
 externalism 272–3; about justification 273; about knowledge 273
- factive attitudes 22  
 faculties, active vs. reactive 99  
 Fairweather, Abrol 293 (n. 8)  
 faith 325–7  
 fallibilism 303–4, 374  
 Feldman, Richard 203 (n. 22), 266 (n. 13), 292 (n. 4)  
 feminist epistemology 305–6, 312, 329 (nn. 7–9)  
 Fish, William 59 (n. 2)  
 Fogelin, Robert 353 (n. 1), 354 (n. 7)  
 foundationalism, epistemological 216; moderate 233; psychological 219–20; strong 234  
 Freud, Sigmund 101 (n. 9)  
 Fricker, Elizabeth 170 (n. 13), 172 (n. 22)  
 Fricker, Miranda 329 (n. 8)  
 Fuller, Steve 329 (n. 11)  
 functional dependence 55
- Galileo 299, 301, 302  
 gambler's fallacy 354 (n. 4)  
 generality problem *see* specification problem  
 genetic dependence 139,  
 Gert, Bernard 331 (n. 28)  
 Gettier, Edmund L. 264 (n. 2)  
 Giaquinto, M. 146 (n. 3)  
 Ginet, Carl 80 (n. 8), 292 (n. 6)  
 Goldberg, Sanford 353 (n. 1)  
 Goldman, Alvin I. 203 (n. 22), 266 (n. 16), 292 (n. 5)  
 Graham, Peter 169 (n. 8)  
 Grayling, A.C. 294 (n. 16)  
 Greco, John 35 (n. 14), 126 (n. 16), 170 (n. 14), 238 (n. 11), 293 (n. 8), 294 (n. 12), 353 (n. 1)  
 Grice, H.P. 125 (n. 15)
- grounding 6–7; causal 6; deductive 299 (*see also* strong foundationalism); epistemic 6; inductive 299 (*see also* moderate foundationalism); internal 273, 280; justificational 6; memorial 8, 62, 71, 76 79 (n. 1); perceptual 16, 219; reliable 255–7; truth-conducive 169, 277  
 grounding conditions 227; vs. possession conditions 227
- Hagaman, Scott 202 (n. 18)  
 hallucination 40; argument from 42; empty 41; memorial 69  
 Hannay, Alastair 101 (n. 6)  
 Hardin, C.L. 124 (n. 10)  
 Hardwig, John 171 (n. 15)  
 Harman, Gilbert 203 (n. 22)  
 Hawthorne, John 202 (n. 15)  
 Hegel, G.W.F. 294 (n. 18)  
 Heil, John 59 (n. 4)  
 Hempel, Carl G. 28 (n. 2)  
 Hopkins, Gerard Manley 323  
 Huemer, Michael xvi, 80 (n. 11), 126 (n. 17), 147 (n. 7), 201 (n. 9), 266 (n. 14, 15, 16), 353 (n. 1), 354 (n. 7), 355 (n. 8, 13)  
 Hume, David 91, 92, 165, 338–9, 340, 342, 354 (n. 3), 359, 373
- idealism 52  
 illusion 40  
 imagination 69, 85 88, 300–1  
 imaging, second-order 88  
 immediacy, epistemic 106; epistemic vs. temporal 106  
 infeasibility *see* defeasibility  
 independence principle 234, 239 (n. 17)  
 indirectness 112; causal 44; epistemic 44, 143; inferential 181; justificational 386; objectual 43, 44; visual 42  
 induction *see* inference, deductive and inductive; intuitive induction; problem of induction  
 inductive reasoning 187; analogical 187; explanatory 187; generalizational 187; vs. perception 179; and probability 189; *see also* inference to the best explanation  
 inductive transmission 192–3  
 infallibility 344–5; thesis of 92  
 inference 80 (n. 7), 152, 153, 177–78; analogical; to the best explanation 300–1; belief-forming 183; confirmatory vs. generative 183; content of vs. process of 177; deductive and inductive 186–7; subsumptive 187; vs. perception 44, 179  
 inferential chains 208  
 inferential dependence 160  
 instantial reasoning 187  
 intellectual courage *see* courage, intellectual  
 intentional object 100 (n. 4)



- internalism 272–3; about justification 273;  
 about knowledge 273  
 intrinsic goodness 282–3  
 intrinsic property 46  
 introspection 3, 85, 87; and perception 88,  
 89–90; as a basis of justification and  
 knowledge 96–100  
 intuition 106, 135, 140, 276, 314,  
 intuitionism (ethical) 317, 318, 330 (n. 15)  
 intuitions 314  
 intuitive induction 126 (n. 22)  
 invariantism 264  
 irrealism, direct 52; *see also* anti-realism  
 isolation problem *see* coherentism
- James, William 294 (n. 19)  
 Jeshion, Robin 122 (n. 2)  
 justification 3, 28–9, 76, 96–8, 158–9, 197,  
 272–5, 367–9; belief justification 3 (*see also*  
 justification, situational); conclusive 250;  
 defeasible *see* defeasibility; dispositional  
 vs. occurrent, 81; degrees of 140, 147  
 (n. 6), 197, 326; doxastic 3; for believing  
*see* justification, situational; global vs.  
 individual 106; in believing *see* justifica-  
 tion, situational; indefeasible 143, 145,  
 236; introspective 97–8; prima facie 28;  
 process of 229; property of 230, 383;  
 second-order 367; situational 3, 8, 25,  
 194, 197, 376 (n. 7); social 322 (*see also*  
 testimony); transmission of 187–97; and  
 truth 281–2; *see also* knowledge  
 justificational dependence 216
- Kant, Immanuel 315, 317, 318, 330 (n. 17)  
 Kirkham, Richard 294 (n. 15)  
 Kitcher, Philip 126 (n. 18)  
 KK thesis 365  
 Klein, Peter D. 203 (n. 22), 237 (nn. 3, 4), 238  
 (n. 9), 265 (n. 8), 354 (n. 7), 355 (n. 13)  
 knowledge: in the abstract 207; approximate  
 304; a priori *see* a priori knowledge; basic  
*see* basic knowledge; by acquaintance vs.  
 description 253; causal account of 254; and  
 certainty 347; clairvoyant 270–2; defeat of  
 183; direct 214, 320; fortunate 260; foun-  
 dational 214; general conception of 291;  
 natural 272; recognitional 155; second-  
 order 337, 365; secondary 317–18; social  
 (*see also* testimony) 306–7; sociology of  
*see* sociology of knowledge; testimonially  
 direct 317; transmission of 187–197; value  
 of *see* value problem; virtual 304; *see also*  
 justification
- Kripke, Saul 146 (n. 4)  
 Kvanvig, Jonathan L. 240 (n. 20), 264 (n. 5),  
 293 (nn. 8, 12)  
 Kyburg, Henry E., Jr. 202 (n. 16)
- Lawrence, D.H. 292 (n. 1)  
 learning 20, 51, 162–3; conceptual vs. propo-  
 sitional 162  
 Lehrer, Keith 168 (n. 3), 239 (n. 15), 240 (n.  
 19)  
 Lennon, Kathleen 329 (n. 7)  
 Lewis, C.I. 60 (n. 9)  
 Lewis, David 266 (n. 15)  
 Leibniz, Gottfried Wilhelm xii, 126 (n. 21),  
 146 (n. 4)  
 Locke, John xii, 43, 44, 64, 79 (n. 5)  
 logical possibility 124 (n. 11), 360  
 logical truth 117–18, 121, 126 (n. 19)  
 Longino, Helen 329 (n. 8)  
 lottery paradox 202 (n. 16)  
 luck 156, 191, 246, 249, 261
- Marcus, Ruth Barcan 123 (n. 7)  
 McDowell, John 170 (n. 13)  
 McGinn, Colin 124 (n. 10)  
 McLaughlin, Brian 101 (n. 7)  
 meaning: change of vs. falsification 134–5;  
 truth by virtue of 130–1  
 memory: adverbial conception of 71; episte-  
 mological centrality of 75; modes of 66;  
 phenomenalist conception of 70; repre-  
 sentative theory of 68–9; *see also* belief;  
 justification; knowledge  
 mentalism 292 (n. 4); *see also* internalism  
 mental properties and states 84–6; *see also*  
 privileged access  
 Mill, John Stuart 116, 117, 126 (n. 18), 315,  
 316–18, 330 (nn. 19, 20)  
 mind-dependence 46, 287  
 Moore, G.E. 147 (n. 5), 265 (n. 8), 355 (n. 13)  
 Moser, Paul K. 239 (n. 16, 354 (n. 7))
- naïve realism *see* realism, naïve  
 naturalistic epistemology 253–4  
 necessary truth 106; *see also* analytic proposi-  
 tion, logical truth  
 necessity, empirical 349  
 necessity, logical *see* logical truth  
 necessity, natural 355 (n. 12)  
 necessity, nomic 136  
 Neta, Ram 293 (n. 12)  
 Newton, Isaac 116, 302, 303, 304  
 nomic necessity *see* necessity, nomic  
 noncognitivism 309  
 Nozick, Robert 203 (n. 22), 266 (n. 14), 354  
 (n. 7)
- objectivism 303  
 observation 254, 328 (nn. 1, 2)  
 occurrent properties 71, 86, 92  
 Odegard, Douglas 203 (n. 22)  
 omniscience, thesis of 92  
 operational dependence 160

- other minds, problem of *see* problem of other minds
- parity, epistemic *see* epistemic parity
- Peirce, C.S. 287
- Pendelbury, Michael 32 (n. 3)
- percept 25
- perceptible 16, 51–2, 90, 176
- perception: aspectual 26; basic elements in 17; and belief 15–24; causal theory of 39; inner and outer 84, 89, 380; and introspection 89–90; moral 200 (n. 3), 313–14; objectual 20; propositional 20; residue of 68; simple 17; vs. inference 44, 179; *see also* adverbial theories; belief, dispositional vs. disposition to believe; realism, naive; phenomenalism; sense-datum theories; theory of appearing
- perceptual hierarchy 24–5
- phenomenal conservatism *see* conservatism, phenomenal
- phenomenal property 87, 100 (n. 3)
- phenomenalism 52; adverbial 52; sense-datum version of 51
- Plantinga, Alvin 170 (n. 9), 264 (nn. 5, 10, 12, 19), 276, 331 (n. 22)
- Plato xi, 79 (n. 3), 117, 246, 264 (n. 1), 282
- possibility, conceptual 124 (n. 11)
- possibility, epistemic vs. logical 360
- possibility, logical 360
- presuppositional dependence 249
- privileged access 91–6; strong doctrine of 92; *see also* infallibility; and omniscience, thesis of
- probability *see* Hume, inductive reasoning, probable reasoning, reason to believe
- probable reasoning 186, 189, 359 *see also* inductive reasoning
- problem of the body 342
- problem of the criterion 355 (n. 13)
- problem of induction 359–61
- problem of other minds 340–2
- problem of too many minds 354 (n. 6)
- proof 299
- property attribution 33 (n. 4)
- Pritchard, Duncan 265 (n. 6), 293 (n. 12)
- property *see* dispositional properties, intrinsic property, occurrent properties; qualities, primary; relational property
- provability 113, 142, 144
- Pryor, James 355 (n. 13)
- psychological dependence 249
- psychological foundationalism *see* foundationalism, psychological
- Putnam, Hilary 353 (n. 1), 355 (n. 8)
- Pyrrho of Ellis 335; *see also* skepticism, Pyrrhonian
- Pyrrhonian skepticism *see* skepticism Pyrrhonian
- qualities, primary 43
- qualities, secondary 43
- Quine, W.V. 125 (n. 15), 145 (n. 1), 146 (nn. 2, 3), 294 (nn. 16, 17), 328 (nn. 1, 2)
- Quine–Duhem thesis *see* Duhem–Quine thesis
- rational disagreement *see* disagreement, rational
- rationalism 116–17
- rationality 325–6
- Reid, Thomas 64, 168 (n. 3), 171 (n. 15)
- realism 38, 329 (n. 14); direct 41, 47; indirect 52, 67; naive 38; perceptual 38, 41; representative 43
- reason 104; to believe xi, 197, 202 (nn. 13, 14), 346, 360, 376 (n. 9); *see also* truths of reason
- reasoning 75, 176, 177, 178, 180, 186; by analogy; probabilistic *see* probabilistic reasoning; second-order, 109; *see also* deductive reasoning; inductive reasoning; inference; subsumptive reasoning
- recalling 66
- regress argument *see* epistemic regress argument
- regress of demonstration 369
- relational property 46
- relativism (ethical) 308–9; circumstantial 312; vs. relativizationism 309
- relevant alternative 203 (n. 20), 261–2, 352; *see also* contrast class
- reliabilism: indicator 265 (n. 11); process 265 (n. 11); *see also* reliability theories
- reliability theories 255–64, 371
- remembering 62–4; active (occurrent) 71; event 66; how 66; objectual 66; propositional 66
- representative realism *see* realism
- Rescher, Nicholas 354 (n. 7)
- Rigg, Wayne D. 240 (n. 20)
- Robinson, Howard 59 (n. 3)
- Rorty, Amélie O. 101 (n. 7)
- Rorty, Richard 200 (n. 8), 239 (n. 18)
- Ross, W.D. 126 (n. 22), 147 (n. 5), 330 (n. 15)
- Russell, Bertrand 126 (nn. 15, 17), 201 (n. 9), 354 (nn. 4, 6)
- Sainsbury, Mark 294 (n. 16)
- Sayre-McCord, Geoffrey 329 (n. 14)
- Schmitt, Frederick 329 (nn. 8, 11)
- scientific knowledge *see* knowledge, scientific; proof; testimony
- seeing 56–8: as a causal relation 17, 27; by 27; and light 56
- seemings 147 (n. 7), 314, 330 (n. 16)
- self-consciousness 97
- self-deception 101 (n. 7)

- self-evidence 105, 112–14, 122 (n. 1); immediate 106; relativized notion of 125 (n. 13), 147 (n. 9)
- self-grounding belief *see* belief, self-grounding
- self-knowledge principle 96
- Sellars, Wilfrid 238 (n. 8), 240 (n. 19)
- Senor, Thomas D. 321 (n. 27)
- sense-data 42; *see also* phenomenalism; sense-datum theories
- sense-datum theories: of non-perceptual sensory experience 50; of perception 41–6
- sensory experience; *see also* adverbial theories; phenomenalism; sense-datum theories
- Shakespeare 48, 56, 86, 179
- Shalkowski, Scott A. 124 (n. 11)
- showing (a proposition) 229; de facto (n. 15); dialectical 356 (n. 15)
- show-know principle 352
- Sinnott-Armstrong, Walter 354 (n. 5), 356 (n. 14)
- situational justification *see* justification, situational
- skeptical principles: back-up 347; certainty 347; cogency 362; entailment 359; infallibility principle proper 345, 361; necessity 344
- skepticism: attitudinal 336; first-order 337; justification skepticism 336; knowledge skepticism 336; moderate vs. radical 366; and the positive defense of common sense 366–71; Pyrrhonian 335, 353 (n. 2); refutation vs. rebuttal of 365; second-order 337
- Smart, J.J.C. 59 (n. 7)
- social epistemology 305–6 *see also* testimony
- sociology of knowledge 312
- Sosa, Ernest 238 (n. 11), 293 (nn. 8, 10, 12), 329 (n. 8), 330 (n. 16), 351 (n. 1), 354 (n. 7), 376 (n. 8)
- source conditions 185; *see also* transmission conditions
- specification problem 258–9, 266 (n. 13)
- Steup, Matthias 147 (n. 7), 293 (n. 12)
- Strawson, P.F. 125 (n. 15)
- Stroud, Barry 353 (n. 1), 354 (n. 7), 376 (n. 8)
- Stump, Eleonore 331 (n. 29)
- subsumptive reasoning 187
- suspended judgment 5
- synthetic a priori 114–17, 124 (n. 10)
- synthetic propositions 112
- Tarski, Alfred 294 (n. 16)
- Taylor, James E. 295 (n. 20)
- testability 118, 313
- testimony: demonstrative vs. propositional 162; formal and informal 150
- theory of appearing 40
- Timmons, Mark C. 35 (n. 14), 294 (n. 13), 330 (n. 14), 353 (n. 1), 356 (n. 14)
- tracking 260, 266 (n. 14)
- transmission conditions 185–97; and testimony 156–7; *see also* source conditions
- trust 152–3
- truth: coherence theory of 288–9; correspondence theory of 286–7; criteria vs. nature of 290; essential 137; for a person or culture 310–11; logical *see* logical truth; minimalist account of 287–8; pragmatic theory of 289–90; redundancy account of 287–8; *see also* justification; knowledge; meaning; necessary truth
- truths of reason 176 *see also* a priori propositions, self-evidence
- uncertainty 347–48; *see also* certainty
- underdetermination 342, 347, 354 (n. 7)
- understanding 104, 105, 111, 121 (n. 1), 133, 141; of and that 111; value of 285
- uniformity of nature principle 359
- vagueness 23, 127 (n. 25), 134, 264
- validity 186
- van Gogh 26
- verity principle 344
- virtue epistemology 277–80
- value problem, the 282–6
- visual experience principle, the 30
- visual justification principle, the 29
- Vogel, Jonathan 354 (n. 7)
- voluntarism, doxastic *see* doxastic
- warrant 170 (n. 10), 264 (n. 5)
- Webb, Mark Owen 171 (n. 15)
- Whitford, Margaret 329 (n. 7)
- Williams, Michael 353 (n. 1), 354 (n. 7)
- Williamson, Timothy 35 (n. 14), 127 (n. 25), 264 (n. 3), 353 (n. 1)
- Wittgenstein, Ludwig 239 (n. 14)
- Yandell, Keith E. 330 (n. 21)
- Zagzebski, Linda 293 (n. 8), 294 (n. 12)